AFIT/GAP/ENP/95S-01

NON-IMAGING INFRARED

SPECTRAL TARGET DETECTION
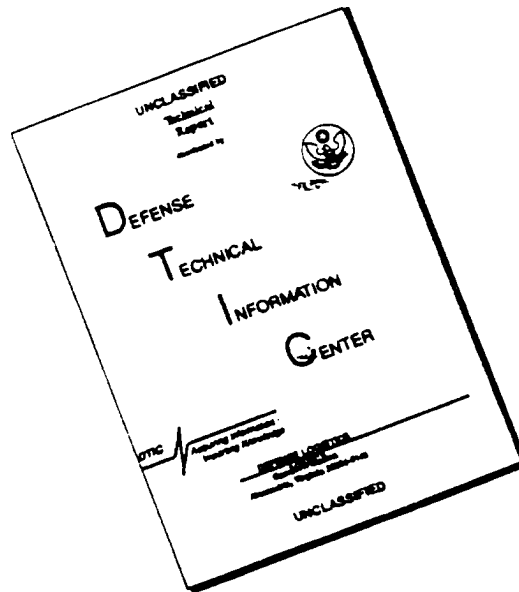
THESIS
Matthew Raymond Whiteley
Captain, USAF

AFIT/GAP/ENP/95S-01

19960327 021

Approved for public release; distribution unlimited

DISCLAIMER NOTICE

THIS DOCUMENT IS BEST
QUALITY AVAILABLE. THE COPY
FURNISHED TO DTIC CONTAINED
A SIGNIFICANT NUMBER OF
PAGES WHICH DO NOT
REPRODUCE LEGIBLY.

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U. S. Government.

NON-IMAGING INFRARED SPECTRAL TARGET DETECTION

THESIS

Presented to the Faculty of the School of Engineering

of the Air Force Institute of Technology

Air University

In Partial Fulfillment of the

Requirements for the Degree of

Master of Science

Matthew Raymond Whiteley, B.S.

Captain, USAF

September, 1995

Acknowledgements

Table of Contents

List of Figures

List of Tables

AFIT/GAP/ENP/95S-01

Abstract

Automatic detection of time-critical mobile targets using spectral-only infrared radiance data is explored. A quantification of the probability of detection, false alarm rate, and total error rate associated with this detection process is provided. A set of classification features is developed for the spectral data, and these features are utilized in a Bayesian classifier singly and in combination to provide target detection. The results of this processing are presented and sensitivity of the class separability to target set, target configuration, diurnal variations, mean contrast, and ambient temperature estimation errors is explored. This work introduces the concept of *atmospheric normalization* of classification features, in which feature values are normalized using an estimate of the ambient temperature surrounding the target being observed and applying the Planck radiation law with those estimates. This technique is demonstrated to reduce the total error rate associated with classification processing to less than one-fourth of that observed using non-normalized features. Classification testing of spectral field measurements made on an array of U.S. and foreign military assets reveal a total error rate near 5% with a 95% probability of detection and a concurrent false alarm rate of 4% when a single classification feature is employed. Multiple feature classification on the same data yields detection probabilities near 97% with a concurrent false alarm rate of 2.5%. Sensitivity analysis indicates that the probability of detection is reduced to 70-75% in the hours preceding daylight, and that for the total error rate to be less than 10%, the target-to-background mean contrast must be greater than 0.1. Analysis of the atmospheric normalization technique reveals that in order to keep the total error rate less than 10%, the ambient temperature must be estimated with less than 3K absolute accuracy.

NON-IMAGING INFRARED SPECTRAL TARGET DETECTION

I. Introduction

1.1 Research Motivation

The results of air campaigns conducted during the Persian Gulf War empha-
sized the relative difficulty of locating and prosecuting time-critical mobile targets
(CMTs). Typical CMTs include ground vehicles such as missile launchers, theater
ballistic missiles, mobile cruise missiles, Scuds, anti-aircraft vehicles, artillery, and
mobile command posts. Since these targets are highly mobile and have the capability
of being relocated within a few hours, wide-area surveillance must be conducted in
order to establish estimates of their locations. Furthermore, since these initial esti-
mates may quickly become outdated, the search process must be repeated frequently.
The requirements for wide-area surveillance and large-area acquisition dictate that
the processes used to locate and prosecute these CMTs must be automated to meet
mission time-lines.

While much research has been accomplished in the area of automated target
detection, recognition, and identification (ATD/R/I), high probabilities of detec-
tion (recognition or identification) and low false alarm rates have often come about
through the use of high spatial resolution remote sensors. High spatial resolution
results in lower area search rates for the sensor and longer image processing time-
lines. These features are inconsistent with the time-lines and search areas required
to locate and prosecute CMTs.

The paradigm of *directed vision* has been proposed to provide for automated
CMT location and prosecution within the relevant time-lines. In this paradigm, a
system with a fast scan rate carries out a coarse classification (e.g. target vs. non-

target). The detection output (declaration of a target detection) is used as a cue to a system that can perform classification at a higher level (e.g. tank vs. truck). The task of the first system is to provide for a high probability of detection with a false alarm rate determined by the ability of the second system to reject these false alarms while keeping up with the cuing rate. The task of the second system is to examine the cues received from the first and perform higher level discrimination functions, thus rejecting false alarms passed to it. The net effect of the two systems working in tandem is to yield high probabilities of detection and low false alarms rates in a time-line that facilitates CMT location and prosecution.

A multispectral infrared sensing system has been proposed as a candidate for use in the directed vision paradigm. With this type of sensor, the target signature is resolved spectrally and the spatial resolution is reduced. With reduced spatial resolution, the area search rate of such a system is consequently increased. At the same time, resolution of the signature in the spectral domain emphasizes material composition of the observed target. The spectral signature is therefore ideally suited for the detection phase in the directed vision paradigm as it applies to the location and prosecution of CMTs.

The ability to accurately distinguish target from non-target using spectral data shows great promise towards increasing area search rates and decreasing processing time-lines. While target detectability has been *assessed* using spectral data, detection capability and associated error rates have yet to be objectively *quantified* on a set of infrared spectral measurements. Without proper quantification, the assumed benefit of a spectral sensor remains speculative.

1.2 Problem Statement

This thesis research focuses on quantifying the probability of target detection, false alarms rate, and total error rate that may be achieved by performing automated classification processing on non-imaging infrared spectral target observation data.

1.3 Current Knowledge

The problem of automating the target detection process and quantifying the performance of that process on a set of measurements falls under the broad category of pattern recognition. As such, the pattern recognition theory and methods that are employed routinely with image data may be applied fully to the non-imaging spectral-only problem. Ultimately, the utility of the automated process that is developed must be tested against a set of measurement data and the results of that processing (classifications) must be compared to desired results (ground truth).

In research currently being carried out by the Wright Laboratory Avionics Directorate (WL/AARI), Naval Research Laboratory (NRL), and the Environmental Research Institute of Michigan (ERIM), the target detection paradigm focuses on exploiting the empirically-observed fact that natural objects or background (tree canopy, grass, etc.) exhibit a high degree of correlation between spectral radiance measurements made at certain wavelength pairs (λ_1, λ_2). Man-made objects or targets (tanks, trucks, etc.) tend to lie farther off the background correlation axis (19). Target and non-target (background) observations are assumed to be normally-distributed when projected onto an axis orthogonal to the background correlation axis, subject to certain requirements on the degree of background correlation. From the observed distributions, a signal-to-clutter ratio (SCR) is measured. Target detectability is then assessed through the classical radar detection model, where an analogy is drawn between the computed SCR and the radar signal-to-noise ratio (SNR). In this detection paradigm, efforts are directed toward finding wavelength pairs producing an SCR commensurate with target detection and false alarm rates that meet the program goals.

In 1994, Thomas (16) reported that for the target set he examined, higher background correlation existed between pairs in the long-wave infrared (LWIR) than in the mid-wave infrared (MWIR) for correlation levels (ρ) of interest, $\rho \geq 0.99$. Additionally, he found that LWIR-MWIR band pairs rarely resulted in high background

correlation. Furthermore, Thomas established a metric which he called SCORE that found more stable band pairs for target detectability estimation than could be obtained by simply calculating SCR values for the same observations. This fact indicates that the observed SCR may not fully parameterize target detectability.

Although the research currently underway has established that targets are detectable in highly-cluttered backgrounds by exploiting background spectral correlation, the detectability must be taken with proper caveat, since the target detection probability and false alarms rates are based upon *modeled* detector behavior as opposed to *measured* detector behavior. The research accomplished in this thesis establishes error rates through the design, implementation, and testing of an alternate detection paradigm and quantifies those rates by objective comparison of detector output with known truth, thus avoiding assumptions made in attempting to model the detection process.

To accomplish the task of quantifying the performance of a non-imaging infrared spectral target detection system, the work draws upon a wide range of radiometric theory established in the past century as well as statistical classification theory and error rate evaluation techniques developed over the past several decades. The key ideas behind this knowledge are presented here, with a more detailed treatment reserved for Chapter II.

1.3.1 Radiometric Theory. The problem of measuring the energy contained in optical radiation fields is the fundamental issue that radiometric theory confronts. Radiometry in its present form deals only with incoherent radiation that follows the laws of geometric optics. In the study of radiometry, certain radiometric quantities have been established. Chief among these radiometric quantities is that of *spectral radiance*, which is defined to be the power emitted from a source per unit area per unit solid angle at a particular wavelength. This quantity is given the symbol L_λ.

4

Of particular interest (both practical and historical) in radiometry is the study of the spectral radiance emitted by a so-called *blackbody* (1). A blackbody is a hypothetical body that absorbs all radiation incident upon it. While a true blackbody is a purely imaginary entity, many common sources behave in a fashion similar to an ideal blackbody. A small hole which is made in an enclosure maintained at a uniform temperature closely approximates the behavior of a blackbody (1:28). Additionally, most physical objects not undergoing other emission processes exhibit emission properties similar to that of a blackbody. The correct theoretical expression for blackbody spectral radiance was discovered by Planck in 1900 (1:51). This distribution, usually referred to as the *Planck radiation law*, is presented and discussed in detail in Chapter II.

1.3.2 Feature Spaces and Statistical Classification. Given a signal (image, spectrum, etc.), we may define and measure some *feature* of that signal. This feature may be either data taken directly from the signal, or it may be derived or calculated from a part of, or the entire signal. Features may assume any form, with the only restriction being that the method of measuring that feature is consistent for all signals of the same type, regardless of the class to which a signal belongs. For example, in this research, the same method must be used to measure a feature value from a spectrum whether that spectrum comes from a target or from a non-target. If we construct N such features (each different), then we define an N-dimensional space into which each of our signals is mapped.

The mapping of signals into feature space is done in order that we may determine to which class of signals the signal of interest belongs. To do this, a set of discriminant functions defined in the feature space must be used. These discriminant functions tell us the probability that the observation in question belongs to a particular class given its location in feature space. Our *classification* of the observation is based upon these probabilities. These concepts are fundamental to the statistical

classification techniques that will be discussed fully in Chapter II and will form the basis of our target detection process.

1.3.3 Error Rate Estimation Techniques. The standard metric for assessing performance in any statistical classification problem is the minimum average probability of error or the *Bayes error rate*. In the target detection problem, the Bayes error rate is composed of both false alarms and missed detections. Given a feature space and a set of observations, a unique value for the Bayes error rate exists as the minimum error rate at which *any* classifier in that feature space may operate. Put another way, the Bayes error rate represents a limit in the separability that may be achieved between classes in a given feature space. Thus, if the Bayes error rate measured for a set of observations exceeds the error that is tolerable, then a new feature space must be developed in which class separability is increased. Changing the type of classifier used will not change the Bayes error rate.

In practice, the Bayes error rate is measured by first estimating the class-conditional probability for each class at each value in the feature space. These values are used in the so-called Bayesian decision rule, with the error resulting from this decision process being taken as an estimate of the Bayes error rate. The issue relevant in whether this error rate is valid lies in the ability to estimate the class-conditional probabilities. This estimate may be made either through parametric or non-parametric techniques. Non-parametric techniques are usually favored over parametric techniques, as they make no assumptions as to the form of the class-conditional distributions. These techniques do, however, require training data that is representative of the data expected during testing. Issues surrounding error rate estimation are discussed in greater depth in Chapter II and Chapter III.

1.4 Spectral Infrared Database

Under the Joint Multispectral Sensor Program (JMSP), which is currently managed by WL/AARI, several data collections of infrared spectral radiance measurements (3-14 μm) have been made. Included in the current JMSP database are measurements conducted at Wright-Patterson AFB in Ohio, Redstone Arsenal in Alabama, White Sands Missile Range in New Mexico, and Camp Grayling in Michigan. In all of these data collection efforts, measurements were made on various target configurations and backgrounds (non-target observations) throughout 24 hour periods over the course of several days or weeks. Each collection of data is organized into 'experiments' which consist of a number of spectra observed from each of several targets and non-targets. Overall, during each of these collection efforts, thousands of target and non-target observations were recorded and were available for use in this research. This data is discussed in greater detail in Chapter III.

1.5 Scope

In this thesis, a set of classification features is developed for the infrared spectral radiance data, and the performance of these features is evaluated both separately and in combination. Class-conditional probability distribution functions (PDFs) are estimated using the JMSP data and the classification feature set. A classifier employing the Bayesian decision rule has been designed and implemented to work with these PDFs. Testing of the classifier is focused on a set of 137 experiments conducted over a five-day period in June 1994 at Wright-Patterson AFB, Ohio consisting of over 2,700 target and 1,500 non-target observations. Detection performance sensitivity to target set, target configuration, observation time, mean contrast, and ambient temperature estimation is quantified. Estimates of the Bayes error rate associated with classification processing are made non-parametrically using resubstitution and leave-one-out testing techniques.

1.6 Approach

The approach taken to assess detection performance for the non-imaging infrared spectral radiance data has several steps. First, a set of classification features is developed. These features are both direct and derived features of the spectral observations. A Bayesian classifier is implemented to work with single features, multiple independent features, and two-dimensional feature spaces in both resubstitution and leave-one-out testing.

With the established feature set and classification algorithm, error rates for the single feature class estimation are quantified and the features are ranked accordingly. Next, performance improvement using multiple features that are assumed to be independent is quantified, focusing on the features from the set that exhibit the top single feature performance. Finally, detection performance is quantified using two-dimensional feature spaces constructed from the top-performing single features. In all cases, the Bayes error rate is estimated using several non-parametric estimation methods.

The sensitivity of the established detection performance to variations in target set, target configuration, and observation time is explored and quantified by considering observations taken with varying targets present, in varying test configurations, and at different times. Observation conditions and the associated detection performance are quantified as a function of mean contrast. Finally, the detection performance sensitivity will be assessed for the top-performing feature as a function of the sensor's ability to determine the ambient radiative temperature at the time of observation.

1.7 Research Objectives

The research objectives to be accomplished in this thesis are as follows:

1. Develop a set of classification features that may be used individually, in combination, or to define a two-dimensional feature space for the spectral infrared data.

2. Design and implement a classifier employing the Bayesian decision rule that makes use of the classification features developed either individually, in combination, or with a two-dimensional feature space.

3. Test the classifier on a set of spectral observations to establish performance levels (quantified by standard metrics) that may be achieved with a single feature, multiple features, or a two-dimensional feature space.

4. Establish the Bayes error rate for each of the tests using non-parametric estimation techniques.

5. Characterize the sensitivity of the feature set and classifier to variations in target set, target configuration, diurnal conditions, mean contrast, and ambient temperature estimation.

1.8 Key Results

The theory and methods presented in this thesis will serve as a foundation upon which the following key results are built:

- Spectral-only data may be used to provide adequately for the detection of time-critical mobile targets, with probabilities of detection of 95% and greater and corresponding false alarm rates between 2.5-4% being demonstrated on field measurements.

- Atmospheric normalization of classification features, a novel invention of this research, is demonstrated to significantly increase class separability between target and non-target observations.

- While high levels of detection performance can be realized in the daylight hours, early morning observations yield lower probabilities of detection (70-75%) as

non-operating targets loose contrast during this period. Obscuration of targets by tree canopy places the targets in a radiant environment where contrast is lowered and similarly reduces detection performance.

- Mean contrast sensitivity measurements reveal that in order for the total error rate to be less than 10%, targets must exhibit a mean contrast of 0.1 or greater as compared to the surrounding background.

1.9 Thesis Organization

The organization of the remainder of this thesis is as follows: Chapter II will discuss the theory pertinent to this research under the topics of radiometry, spectral correlation, and statistical classification. Chapter III will describe the methodology used throughout the research, offering a characterization of the testing data, definition and physical interpretation of the classification features, description of the class-conditional distribution function generation, classification testing, Bayes error rate estimation, and standardized performance metrics. Chapter IV will be used to present and analyze the results obtained in this research, and those results will be summarized and conclusions will be drawn from them in Chapter V.

1.10 Summary

In order to assess the capability of a non-imaging spectral infrared system in the directed vision concept to detect time-critical mobile targets, detection processing must be carried out and the results of this processing must be compared to the desired output to characterize performance in terms of standardized metrics, without invoking modeled behavior in so doing. The feature set, classification algorithm, and software developed here will permit this detection processing to be carried out on infrared spectral data collected as part of the Joint Multispectral Sensor Program. The Bayes error rate of the established feature set and testing data will be quantified.

Capability will be further assessed by means of sensitivity analysis, exploring the effect of several variables upon a set of standardized performance metrics.

II. Theory

2.1 Introduction

The work accomplished in this thesis is well-grounded in a wide body of knowledge that has come to light over the past century. This theory falls into two broad areas: radiometry and statistical classification. While these areas are not formally related and possess no deep-rooted interconnections, the synthesis of this research effort lies in the intersection of the two fields. Thus, in this discussion, focus is given only to the particular aspects of each area that pertain directly to the problem at hand.

- From the field of radiometry, the following topics are discussed:

    1. pertinent radiometric quantities; their meaning and units of measurement

    2. blackbody radiation; its character and analytic form

- From the field of statistical classification, the following topics are discussed:

    1. feature spaces and discriminant functions

    2. probability and Bayesian decision theory

    3. non-parametric Bayes error estimation

2.2 Radiometry

As mentioned in section 1.3.1, radiometry concerns itself with measuring the energy content of optical radiation fields and determining how this energy flows through optical systems (1:13). Although optical radiation can be classified as fully coherent, partially coherent, or incoherent, radiometry usually treats only incoherent radiation. Treatment of fully coherent radiation requires knowledge of the field components and the associated Poynting vector. For this work, however, knowledge of the incoherent properties of optical radiation will suffice, as this research

deals specifically with the radiometry of thermal sources, which may be viewed as emitting essentially incoherent radiation. By calling this radiation incoherent, it is assumed that the amplitude and phase of its individual radiators vary randomly and independently in space.

2.2.1 Radiometric Quantities. In order to discuss the radiometric properties of a source, it is necessary to first define the properties of interest. Although the definitions are somewhat arbitrary, the convention of Boyd (1) is followed throughout, as his convention is based upon common usage by those who work in the field. While a large number of these radiometric quantities may be defined, only those used in this research are presented here.

The total energy contained in a radiation field is called the radiant energy and is given the symbol Q, having units of joules (J). The time rate of change of this energy dQ/dt is referred to as the radiant flux Φ, which has the units of J/s or watts (W). With this definition, we may thus define another quantity called radiance, L, as the flux per unit projected area per unit solid angle leaving a source. Radiance has units of $W/m^2 \cdot sr$. Mathematically, radiance is defined as:

$$L = \frac{d^2\Phi}{dA_{proj}d\Omega} \tag{1}$$

While the measure of radiance as defined above may be adequate for some sources, it is often instructive to consider the variation of this quantity as a function of frequency or wavelength, its spectral composition. Thus, it is necessary to work with *spectral radiance* L_λ, which is defined as the radiant flux per unit projected area per unit solid angle per unit wavelength and has units of $W/m^2 \cdot sr \cdot m$ for wavelengths measured in meters. Mathematically, spectral radiance is defined as:

$$L_\lambda = \frac{d^2\Phi}{dA_{proj}d\Omega d\lambda} \tag{2}$$

13

Radiance and spectral radiance are therefore related simply by:

$$L = \int_0^\infty L_\lambda d\lambda \qquad (3)$$

In common practice, spectral radiance is often expressed in units other than those cited above, as typical fluxes are more easily measured in micro-watts, receiver areas are measured in square centimeters, and wavelengths are measured in microns. Thus, spectral radiance is often reported in units of so-called micro-flicks ($\mu flick$) which are defined as:

$$1\mu flick = \frac{\mu W}{cm^2 \cdot str \cdot \mu m} = 10^4 \frac{W}{m^2 \cdot str \cdot m} \qquad (4)$$

In this thesis, the $\mu flick$ will be adopted as the standard unit for measurements of spectral radiance.

2.2.2 Blackbody Radiation. As discussed in section 1.3.1, a blackbody is a hypothetical body that absorbs all radiation incident upon it, thus, it is black in the normal sense of the word. While the concept of a blackbody represents an ideal, in practice, an approximation to a blackbody may be formed by making a small hole in the wall of an isothermal cavity. The radiation which exits this type of cavity resembles that of an ideal blackbody (1:28). Furthermore, many objects not emitting or reflecting radiation by other processes tend to behave in a manner similar to that of an ideal blackbody.

Generally, determining directional and spectral properties of the radiation emitted from a thermal source is the intent of radiometry. While this is a complex problem for an arbitrary source, it is much simplified for an ideal blackbody. An analytical description of the spectral properties of blackbody radiation was correctly established by Planck in 1900 (1:39). It is for that reason that the expression for blackbody spectral radiance is often referred to as the Planck radiation law.

14
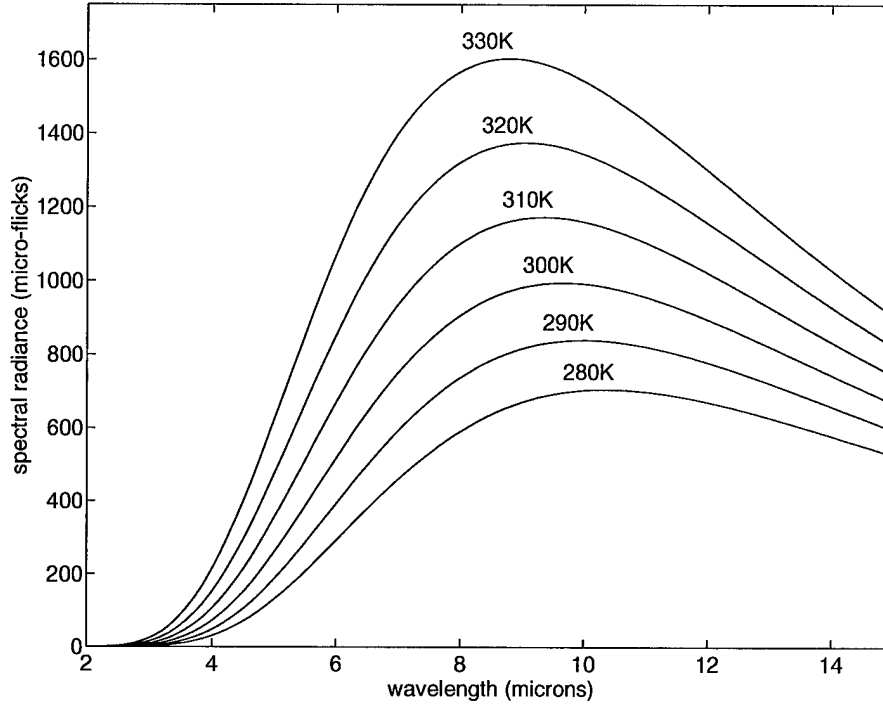
Figure 1. Blackbody spectral radiance for sources at T=280K to T=330K. These curves do not intersect at any wavelength. Thus, at a particular wavelength, the temperature of a blackbody may be uniquely determined.

The Planck radiation law for blackbody sources yields the following form for the spectral radiance of such a source (1:54):

$$L_\lambda^{BB} = \frac{2hc^2}{\lambda^5} \frac{1}{e^{hc/\lambda kT} - 1} \tag{5}$$

where Planck's constant $h = 6.6261 \times 10^{-34} Js$, the speed of light $c = 2.9979 \times 10^8 m/s$, Boltzmann's constant $k = 1.3807 \times 10^{-23} J/K$, and T is the temperature measured in Kelvin. For wavelengths from 2-15 microns and temperatures from 280-330K at 10K increments, the set of curves depicted in figure 1 result from this analytic form of blackbody spectral radiance.

15

As can be seen from examining the form of equation 5 and is obvious from the curves in figure 1, the spectral radiance of a blackbody can be completely specified in terms of a single parameter T, temperature. Additionally, at no point do these curves intersect. Thus, at a particular wavelength, the temperature of a blackbody may be uniquely identified by solving equation 5 for T. This gives us an expression for the so-called *radiant temperature* T_R (1:60):

$$T_R = \frac{hc}{\lambda k \ln\left[1 + 2hc^2/\lambda^5 L_\lambda(\lambda)\right]}.$$
(6)

The spectral radiance emitted from a particular source will in general be different than that of a blackbody at the same temperature. The ratio of the spectral radiance of a source to that of a blackbody under the same geometry is called the *spectral goniometric emissivity* ε of that source and is defined as (1:61):

$$\varepsilon(\theta, \phi, \lambda, T) = \frac{L_\lambda(\theta, \phi, \lambda, T)|_{source}}{L_\lambda(\theta, \phi, \lambda, T)|_{blackbody}}.$$
(7)

2.2.3 Spectral Correlation. As mentioned in Chapter I, much work has been done in exploiting spectral correlation as a detection feature for target spectral observations. Using the analytic form of blackbody spectral radiance, it may be shown that a blackbody is perfectly correlated spectrally. If a wavelength pair (λ_1, λ_2) is chosen, then a blackbody at temperature T may be located in the spectral radiance correlation plane at the coordinate $(L_\lambda(\lambda_1, T), L_\lambda(\lambda_2, T))$. If a plot is made of these coordinates for several values of T, those points will lie along a line, as is shown in figure 2.

In general, the curve given in parametric form as $[L_\lambda(\lambda_1, T), L_\lambda(\lambda_2, T)]$ does not specify $L_\lambda(\lambda_2)$ as a linear function of $L_\lambda(\lambda_1)$. For temperatures near 300K however, this curve is sufficiently modeled as being a straight line. Furthermore, the slope and intercept of this line are completely determined by the particular wavelength pair (λ_1, λ_2) that is chosen. This will be referred to as the *blackbody correlation line.*

16
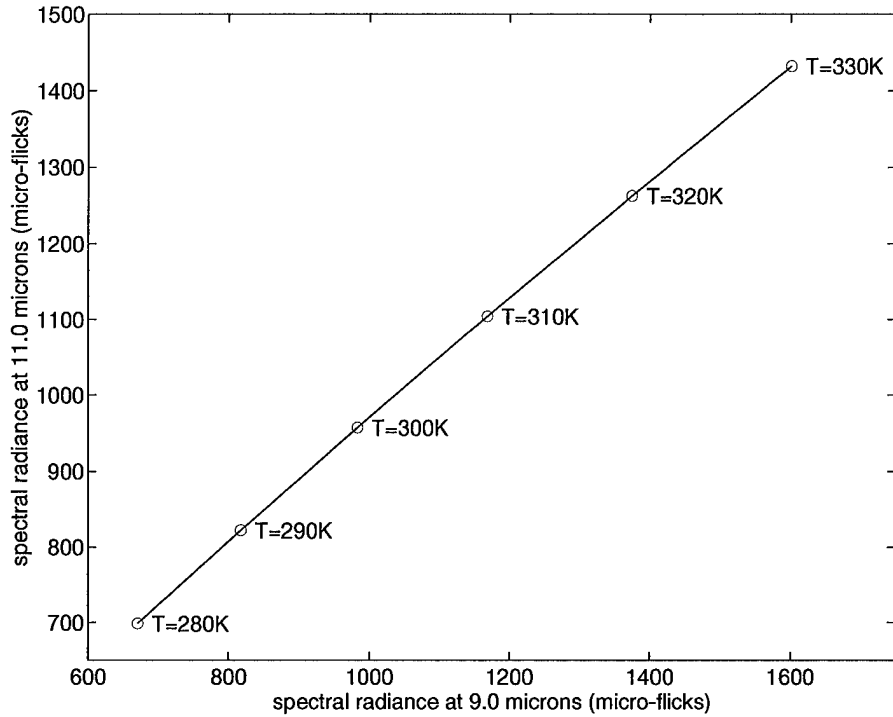
Figure 2. Blackbody spectral correlation for band pair $(9\mu m, 11\mu m)$. This plot indicates the location of a blackbody at several temperatures in the range 280K to 330K.

Spectral observations that lie along the blackbody correlation line exhibit spectral radiance properties that are similar to a blackbody at the particular choice of wavelengths (λ_1, λ_2). Deviations from this line are due to a combination of sensor noise and spectral emissivity variations. In the absence of significant sensor noise, spectral emissivity variations will determine how far a particular spectral observation will lie from the blackbody correlation line. Thus, the empirically-based observation that natural objects tend to be more highly correlated at certain wavelengths than man-made objects in the spectral radiance plane (16:15) is a result of the spectral emissivity properties of natural objects being more similar to those of an ideal blackbody.

17

2.2.4 Atmospheric Absorption. When field measurements of spectral radiance are made, the radiation being emitted from a source is subject to absorption processes in the atmosphere along its path to the receiver. These processes attenuate the signal that reaches the receiver. This attenuation, however, is not constant with wavelength, but instead exhibits a wide range of variation due to the molecular absorption properties of the constituent atmospheric gases such as H_2O and CO_2. The combined effect of such gases in the infrared spectrum is to produce regions of low (atmospheric windows) and high (atmospheric doors) absorption.

The problem of atmospheric absorption is a difficult one, as there are a large number of variables affecting the magnitude of that absorption (i.e., atmospheric composition, constituent concentration, and radiation path). For the purposes of this research, exact quantification of absorption coefficients are not necessary. Instead, it is important to understand in what regions of the spectrum the radiance measurements made are significantly affected, and are therefore less representative of the spectral radiance properties of the observed source.

Figure 3, taken from the IR/EO Handbook (15:40) shows the spectral absorption properties of various gases present in the atmosphere. In general, the observed spectrum will be subject to processes which produce spectral absorption properties that are a combination of the individual gas variations depicted in figure 3. As indicated by the absorption curves, atmospheric windows exist in the 3.0-5.5 μm and 7.3-13.5 μm region of the spectrum. Within the 3.0-5.5 μm window, however, there is a strong absorption line centered near 4.3 μm due to atmospheric CO_2. In the 5.5-7.3 μm spectral region, an H_2O absorption band exists, with absorption in that region being near 100%.

2.3 Statistical Classification

In order to accomplish the task of automated detection of targets from infrared spectral data, it will be required to classify all observations into one of two classes;

18
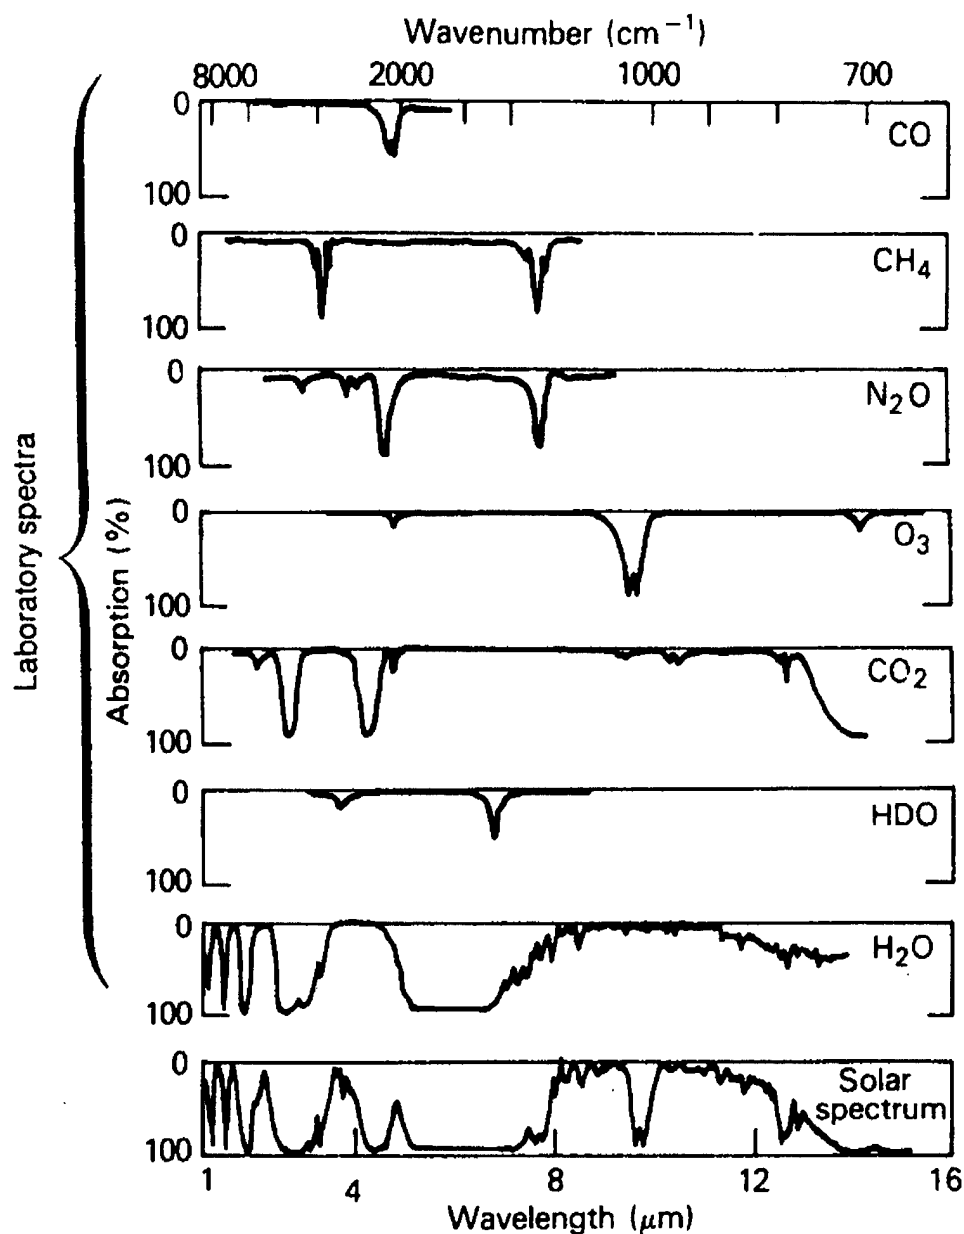
Figure 3. Laboratory absorption spectra of atmospheric gases. Bottom graph indicates combined effect of atmospheric gases on the solar emission spectrum. (Figure reproduced from *The Infrared and Electro-Optical Systems Handbook*, volume 2, page 40, figure 1.13 (15).)

targets and *non-targets*. Thus, the classification problem at hand is a two-class problem, where the decision to be made is to which of the two classes a certain observation belongs. The act of classifying any signal is ultimately little more than an intelligently-formed guess at class membership based upon evidence that is gathered from that signal. To make this guess as intelligent as possible, it is necessary to invoke the mathematics of *statistical decision theory*. The discussion that follows introduces the concepts that are key in this theory and provides guidance for their implementation.

2.3.1 Feature Spaces and Discriminant Functions. In our common experience, classification is a familiar process. We classify the elements of the observed world typically in an unconscious fashion. When asked to characterize the members of a particular class, we often invoke qualitative descriptions that highlight the features of that class and distinguish it from other classes. In this way, we provide for class separation on the basis of observed features.

Quantitatively, we may perform a similar act of class separation of signals (images, spectra, etc.) by the measurement of quantitative features. In a general sense, a feature is any measurement that may be performed on a signal. In the case of images, this may be the value of a gray-level at a scene location. In the case of spectral data, a feature may be, for instance, the spectral radiance at a particular wavelength. Derivative features may also be used whose value is the result of performing a calculation on the signal as opposed to extracting information directly. Whatever the feature is, it must be measured in the same manner for all signals, despite class membership.

With a single feature, or a set of features, measurements made on a signal will yield a vector x that describes the location of that signal in a space defined by the feature or combination of features. It is the location given by x that provides a means by which a signal may be assigned a class. The act of classification or

20

class estimation is performed via *discriminant functions* (3, 4, 6, 14) defined in the feature space constructed from the set of features, X. The argument of these discriminant functions is the feature vector x, the values of which are measured from the signal. These functions are symbolized as $D_i(x)$, where the index i refers to a particular class. For the two-class problem, $i = 1$ (target) or $i = 2$ (non-target). With the use of these discriminant functions, a decision rule may then be defined such as: if $D_1(x) \geq D_2(x)$, then the signal associated with x belongs to class 1 (target), otherwise, it belongs to class 2 (non-target). Thus, with this decision rule, all observations may be assigned a class.

2.3.2 Probability and Bayesian Decision Rule. In the preceding discussion, the concept of a discriminant function defined in a feature space was introduced as a means by which a classification decision may be made. The particular form or origin of these discriminant functions was not discussed. In general, discriminant functions may have any form. It is desirable, however, to use discriminant functions which provide for maximum separation of the classes and minimal total decision error.

To achieve the goal of minimum total decision error for a set of observations, knowledge of the statistical properties of the feature values x for that set must be used. In addition, it is desirable to also have knowledge of the statistical properties of class membership. Specifically, estimates of the class-conditional probability, $p_x(x|\theta_i)$ and the class *a priori* probability (12:18), $P(\theta_i)$ are needed. The class-conditional probability (12:13), $p_x(x|\theta_i)$ (read 'the probability of x given θ_i') is the probability that an observation will have a feature value x if it is of class θ_i. The a priori probability $P(\theta_i)$ is the probability that any observation will belong to the class θ_i. Here, the term 'probability' refers to the relative frequency of occurrence (12:12). For example, if there are a total of 10 balls, 3 red, 5 blue, and 2 green, then $P(red) = \frac{3}{10} = 0.3$, $P(blue) = \frac{5}{10} = 0.5$, $P(green) = \frac{2}{10} = 0.2$.

In the classification problem at hand, a feature value vector x will be measured, and from that measurement a class must be assigned to the observation. Thus, the relevant probability in this process is $p_i(\theta_i|x)$ (read 'the probability of θ_i given x'). This is called the *a posteriori* probability (12:18). This probability may be calculated using Bayes Theorem (12:16):

$$p_i(\theta_i|x) = \frac{p_x(x|\theta_i)P(\theta_i)}{p(x)} \tag{8}$$

The a posteriori probability given by Bayes Theorem in equation (8) may then be used as a discriminant function $D_i(x)$ for the class θ_i. Hence, for the two class problem:

$$\begin{aligned} D_1(x) &= p_1(\theta_1|x) \\ D_2(x) &= p_2(\theta_2|x) \end{aligned} \tag{9}$$

where $\theta_1 = target$ and $\theta_2 = non\text{-}target$. Thus, the a posteriori probabilities may also be used to formulate a decision rule:

$$\begin{aligned} if \ p_1(\theta_1|x) &\geq p_2(\theta_2|x) \ then \ class \ = \ \theta_1 \\ if \ p_2(\theta_1|x) &< p_2(\theta_2|x) \ then \ class \ = \ \theta_2 \end{aligned} \tag{10}$$

This rule is referred to as the *Bayesian decision rule* or the *maximum a posteriori* (MAP) classification criterion, as the class with the maximum a posteriori probability is assigned to the observation with feature measurement vector x.

The implementation of the Bayesian decision rule for classification, in accordance with equation (8) requires knowledge of $p_x(x|\theta_i)$, $P(\theta_i)$, and $p(x)$. The value of $p(x)$ is the same regardless of class and thus its appearance in the a posteriori probability (and therefore in the discriminant functions) is inconsequential in the Bayesian decision rule. Furthermore, in the absence of information to the contrary,

if the occurrence of a target or non-target in an observation is equally likely, then $P(\theta_1) = P(\theta_2) = 0.5$. Hence, in this case, the a priori probabilities are also inconsequential in the decision rule. With these simplifications in hand, the problem of comparing the class a posteriori probabilities only requires estimating the class-conditional probabilities $p_x(x|\theta_i)$, as only a relative comparison of the a posteriori probabilities is required by the Bayesian decision rule. The probabilities $p_x(x|\theta_i)$ are often referred to as the class-conditional probability distribution functions or PDFs. The method of generation of these PDFs is discussed in Chapter III.

2.3.3 Non-parametric Bayes Error Estimation. In the two-class classification problem, two types of errors may occur; a target may be classified as a non-target (missed detection) or a non-target may be classified as a target (false alarm). Both types of error contribute to the total error associated with the classification process. For the Bayesian decision rule with the a priori probabilities assumed equal, these errors occur when the feature vector x_0 computed for an observation lies in a region where the class-conditional probability of the opposite class is greater than the class-conditional probability of the observation's actual class. Graphically, this corresponds to the area of overlap of the two PDFs in feature space, as pictured in figure 4. This error is referred to as the *Bayes error rate* and it represents the minimum average probability of error in classification (4, 3, 6).

The Bayes error rate for a set of observations in a particular feature space is determined by the form of the class-conditional PDFs utilized in the classification process. To construct these PDFs, a set of representative observations called 'training data' is employed. Statistical properties of this data set, such as the mean and standard deviation of a particular feature for each class may be calculated. The PDFs may then be constructed by assuming a functional form, such as a Gaussian distribution, using the characteristic measures of each class as parameters for the distribution. This represents a *parametric* method of PDF construction and the corresponding Bayes error is considered to be a parametric estimation.
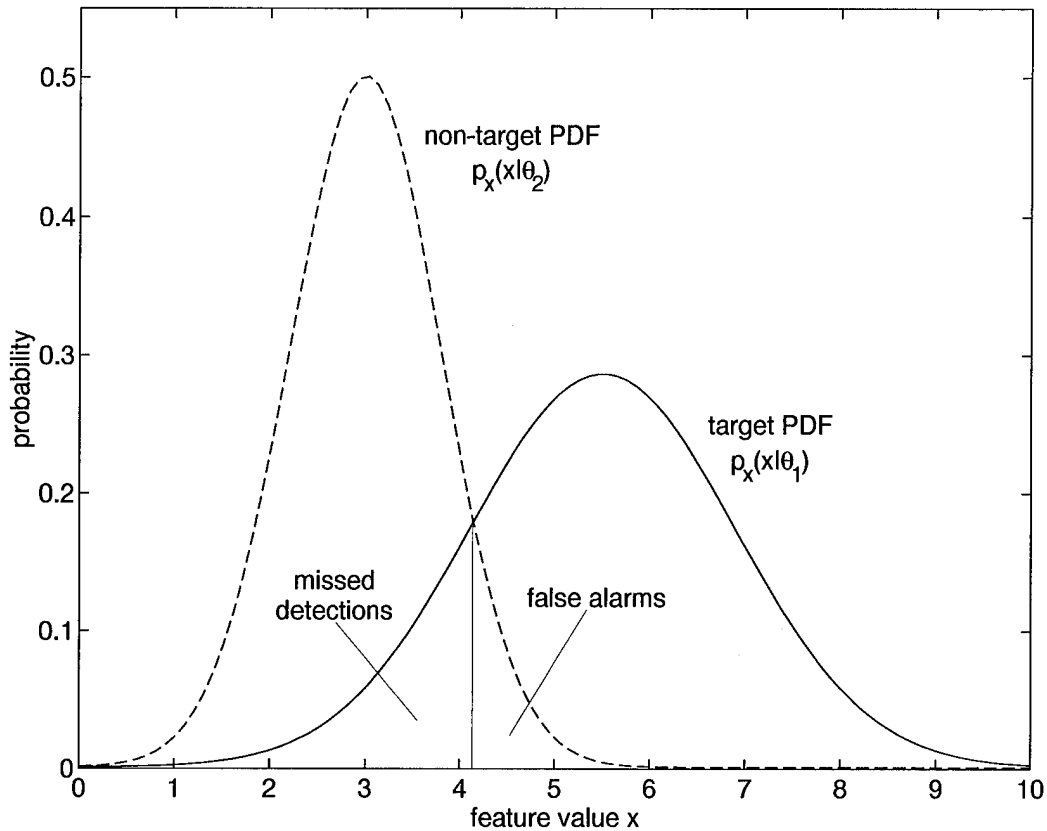
Figure 4. Composition of total decision error.

If a form of the class-conditional PDFs is not assumed, but rather measured from the training set data directly, then the resulting form of the PDFs is *non-parametric*. Correspondingly, a non-parametric Bayes error rate estimate may be made. Non-parametric methods are generally preferred, as the class-conditional PDFs for real observation data rarely follow the exact analytic forms assumed in parametric methods.

Non-parametric construction of class-conditional PDFs may be carried out in several manners. In this research, three density estimation techniques are employed: discrete bins, k-nearest neighbors, and Parzen windows. Each technique is discussed briefly here.

2.3.3.1 Discrete-Bin Density Estimation. The simplest technique for non-parametric class-conditional PDF construction is the use of a histogram (3:424). This involves dividing the feature space into N equally-sized bins. Each bin represents a discrete region in feature space. The probability value for each bin may be calculated based on the relative frequency of occurrence of an observation in that bin (18:40):

$$p_x(x(j)|\theta_i) = \frac{n(x(j), \theta_i)}{n(\theta_i)} \qquad (11)$$

where $p_x(x(j)|\theta_i)$ is the discrete class-conditional PDF value for class i, $x(j)$ represents the feature value x being located in the j^{th} bin, $n(x(j), \theta_i)$ is the number of occurrences of the feature x in bin j for class i, and $n(\theta_i)$ is the number of occurences of class i in the training set. The range of j is from 1 to N.

2.3.3.2 k-Nearest Neighbor Density Estimation. The k-nearest neighbor (k-NN) PDF estimation technique was developed by Fix and Hodges in 1951 and is now used as a standard technique for non-parametric PDF estimation (5). If the volume about a point x in feature space that encompasses the k^{th} nearest neighbors to the point x is small, then the probability estimate at the point x should be large. Conversely, if this volume is large, the probability at x should be small. This idea is the basis for the k-NN PDF estimation technique. Mathematically, this density estimate is expressed as (10:15):

$$p_x(x|\theta_i) = \frac{k-1}{N_i V_d |\Sigma_i|^{1/2} \sqrt{\left[d_i^2(x, x_{k-NN}^{(i)}) \right]^n}} \qquad (12)$$

where N_i is the number of observations of class θ_i, V_d is the volume of an n-dimensional unit radius hypersphere which for n an even integer is (4:24):

$$V_d = \frac{\pi^{n/2}}{(n/2)!},$$

and for n an odd integer is given by:

$$V_d = \frac{2^n \pi^{(n-1)/2} \left(\frac{n-1}{2}\right)!}{n!},$$

Σ_i is the covariance matrix for the i^{th} class (12:127):

$$\Sigma_i = E\left[\left(x^{(i)} - \bar{X}^{(i)}\right)\left(x^{(i)} - \bar{X}^{(i)}\right)^T\right],$$

where the feature vectors $x^{(i)}$ compose the set $X^{(i)}$ for class i, and $d_i^2(x,y)$ is the squared Mahalanobis distance metric for the i^{th} class given by given by (4:38):

$$d_i^2(x,y) = (x - y)^T \Sigma_i^{-1}(x - y).$$

2.3.3.3 Parzen Window Density Estimation.

Parzen introduced an alternative non-parametric PDF estimation technique in 1961 (11). The Parzen window density estimate may be thought of as the sum of windowing functions evaluated at a point x in feature space and centered at an observation point $x^{(i)}$ associated with the class i. The result of this summing is to interpolate between observation points $x^{(i)}$ where each observation contributes to the estimate according to its distance from the point x.

The Parzen density estimate for the PDF of class θ_i in a feature space of dimension n at a feature value x is written mathematically as (7:634):

$$p_x(x|\theta_i) = \frac{1}{N_i}\sum_{j=1}^{N_i}\frac{1}{h^n}k_i\left(\frac{x - x_j^{(i)}}{h}\right),$$

where $x_j^{(i)}$ represents N_i observations of class θ_k, $k_i(\cdot)$ is the window function (kernel) whose spread is controlled by the parameter h. Although the choice of kernel for the Parzen window technique is arbitrary, it is common to use the Gaussian (normal) function. When this kernel is inserted in the general expression for the Parzen density

26

estimate, the expression obtained is (10:16):

$$p_x(x|\theta_i) = \frac{1}{N_i(2\pi)^n/2|\Sigma_i|^{1/2}h^{n+1}} \sum_{j=1}^{N_i} \exp\left(-\frac{d_i^2(x, x_j^{(i)})}{2h^2}\right), \qquad (13)$$

where Σ_i and $d_i^2(\cdot, \cdot)$ are the covariance matrix and squared Mahalanobis distance metric respectively, as defined in section 2.3.3.2.

2.3.4 Classifier Testing. To design and test a classifier and quantify its performance in terms of the Bayes error rate, it is desirable to have a large training set that is separate from the testing set and to be able to completely specify the class-conditional and the a-priori probabilities for each class. Unfortunately, practical data rarely supports this need. To approach the problem of small data sample size, the *resubstitution* and *leave-one-out* testing techniques have been developed. These techniques are also significant since the actual Bayes error rate lies in an interval bounded below by the resubstitution estimate and above by the leave-one-out estimate (6). These techniques are discussed in the sections that follow.

2.3.4.1 Resubstitution. In the resubstitution technique, the observation set is not partitioned into a training set and testing set. Instead, the entire set is used both for training and for testing. The probability of error (portion of the observations misclassified) is assumed to be representative of the error that would be obtained when testing against a similar set of data in the future. While this technique provides a large set of data for training, the resulting error is optimistically-biased, as the class-conditional PDFs developed in training also represent the exact statistical properties of the testing set. Realistically, even similar data obtained in the future is likely not to have the exact statistical properties as the training set. Nonetheless, the resubstitution error rate can still be viewed as the lower-bound on the Bayes error rate, as it represents the 'best-case' performance that may be achieved (6:220).

2.3.4.2 Leave-One-Out. Leave-one-out testing is usually credited to Lachenbruch (9). This type of testing is performed by first taking one observation out of the complete set of observations available. The class-conditional PDFs are then constructed (through any technique) using the remaining available observations. Next, the observation that was taken out is then classified using the PDFs generated from the remaining set. That observation is then placed back into the complete set, another observation is taken out, and the process is repeated. This technique has the advantage of making maximum use of the available observations for training purposes, without optimistically biasing the final classification results. In fact, the error rate generated by leave-one-out testing forms an upper-bound on the Bayes error rate (6:221).

2.4 Summary

This chapter has presented a range of theoretical concepts that are employed in the research presented in this thesis. The key points in this chapter are presented below:

- Optical radiation may be measured in quantities such as spectral radiance which describe the variation of optical field flux density with the wavelength of radiation. The unit of spectral radiance employed in this research is the $\mu flick$, which is a μW per cm^2 per *steradian* per μm.

- The spectral radiance emitted from a blackbody follows the Planck radiation law. This law indicates that blackbody spectral radiance is completely specified in terms of the parameter T, the source temperature. Variations from an ideal blackbody's spectral radiance define a material's spectral emissivity. These variations result in reduced correlation when the spectral radiance of a source at two wavelengths is plotted in a plane. In such a plane, a blackbody therefore exhibits perfect correlation.

- Spectral observations may be treated like other signals and thus, statistical classification theory may be applied directly to the problem of infrared spectral target detection. In order to do this, a feature set must first be established that provides for class separation. With this established feature set, discriminant functions based upon Bayes Theorem may be constructed and invoked in the Bayesian decision rule in order to carry out classification of spectral observations.

- The minimal average error rate associated with classifying a set of observations in a particular feature space is called the Bayes error rate. This rate may be estimated using one or more of several available techniques for constructing the form of the discriminant functions used in the Bayesian decision rule. The bounding of this error rate is accomplished using resubstitution and leave-one-out classifier testing techniques.

In the next chapter, the theory outlined above will be reconciled with the infrared spectral data available. Methodology for the implementation and testing of a target detection algorithm is presented in detail.

III. Methodology

3.1 Introduction

The theory presented in Chapter II provides a strong basis upon which a classifier may be constructed in order to quantify target detection performance with spectral-only infrared data. The present chapter discusses the methods by which this theory was put to use in classifying the JMSP radiance data obtained from the research sponsor, WL/AARI. This chapter is intended to offer the interested reader insight as to how the results presented and analyzed in Chapter IV were obtained. The topics discussed are as follows:

- Testing data: characteristics and specific aspects of the data collection

- Feature set development: techniques developed to provide class separability and definition of established classification feature set

- Class-conditional PDF generation: standard methods and options utilized

- Multiple feature classification: methods for combining feature information

- Classification testing methods: standards for performing tests of the classifier

- Bayes error rate estimation: methods and code used to establish error rates

- Standardized performance metrics: definition of performance quantification

These methods are implemented in a software system called the Infrared Spectra Target Detector, which was designed and developed by the author using the MATLAB numeric computation and visualization system. For details concerning the operation of the detector software, see Appendix A.

3.2 Testing Data

As discussed in section 1.4, the spectral data available for classifier design and testing came from the Joint Multispectral Sensor Program (JMSP) and was

30

made available for this research by WL/AARI. At this point, data from 7 different field collection efforts are available in the JMSP database. Taken as a whole, this database represents well over 10,000 spectral observations taken on approximately 20 different targets along with spectral observations on background material such as tree canopy, grass, soil, gravel, and concrete, including seasonal and diurnal variations. Characteristics of this data are discussed below along with details on the collection efforts.

3.2.1 Spectral Data Characteristics. The spectral data used in this research was collected by the Bomem MB-100 Fourier Transform Spectrometer (FTS) (2), which has been specifically designed by Bomem for radiometric measurements and field use. The MB-100 operates by scanning out an interferogram with its wishbone scan arm on which two corner-cube reflectors are mounted that each receive part of the input radiation from a KBr beamsplitter (2). The interferogram is sampled in conjunction with a reference HeNe laser beam and then Fourier transformed in order to produce the spectrum of the input radiation.

In the configuration used for data collection, the interferometer has two inputs; one directed toward the target of interest and the other to a stable liquid nitrogen reference. There are also two outputs which direct the incident radiation onto two separate detectors; an InSb detector for the MWIR (3-5 microns) and an HgCdTe detector for the LWIR (8-12 microns). An interferogram is produced on each of these detectors (2).

A 10-inch Cassegrain telescope and collimator serve as input optics for the Bomem FTS. The system has a maximum instantaneous field-of-view of 5 milliradians (when limiting apertures are fully opened). A CCD camera is boresighted through the telescope aperture to provide video monitoring of the targets on which spectral radiance measurements are being made. The entire system assembly is mounted on a motorized altitude-azimuth mount which is computer controlled. Each

collection event, referred to as an 'experiment' is completely automated. Once programmed, it is typical for 50 or more spectra to be measured in each experiment, taking 10-15 minutes to collect the entire set.

In a sensor characterization exercise conducted and reported by ERIM in May, 1994 (2), the sensor was found to exhibit exceptional performance. In the tests conducted, a noise equivalent spectral radiance of $10^{-7} \mu flicks$ in the MWIR and $6 \times 10^{-7} \mu flicks$ in the LWIR was measured. Furthermore, the sensor was found to have a relative calibration accuracy of less than 0.25% above 4 μm and spectral correlation preservation greater than 0.9999 in the LWIR. Between the MWIR and LWIR as well as within the 4 to 5 μm region of the MWIR, the correlation preservation was measured to be greater than 0.9995 (2). Since the field measurements made have spectral radiance values on the order of 10^2 to 10^3 $\mu flicks$ for the spectral regions of interest, sensor noise is regarded as being insignificant and its effects are not considered in this research.

Each spectrum that is measured by the sensor is stored in a linear array with 728 elements. This array represents spectral radiance measurements (reported in $\mu flicks$) made between 698.1 and 3502.1 wavenumbers (cm^{-1}). These values correspond to wavelengths ranging from 14.32 to 2.86 μm respectively. Radiance measurements are evenly spaced in wavenumber with a spectral resolution of approximately 4 cm^{-1}. A typical target spectral measurement is shown in figure 5.

Arrays for each measurement in an experiment (typically over 50) are pieced together end-to-end and recorded in one binary file. Ground truth files for each experiment identify the spectra and specify their location in the composite array structure. Associated header files provide collection date and time-of-day information for each experiment. Extensive use of the ground truth and header file information is made by the Infrared Spectral Target Detector software (see Appendix A).
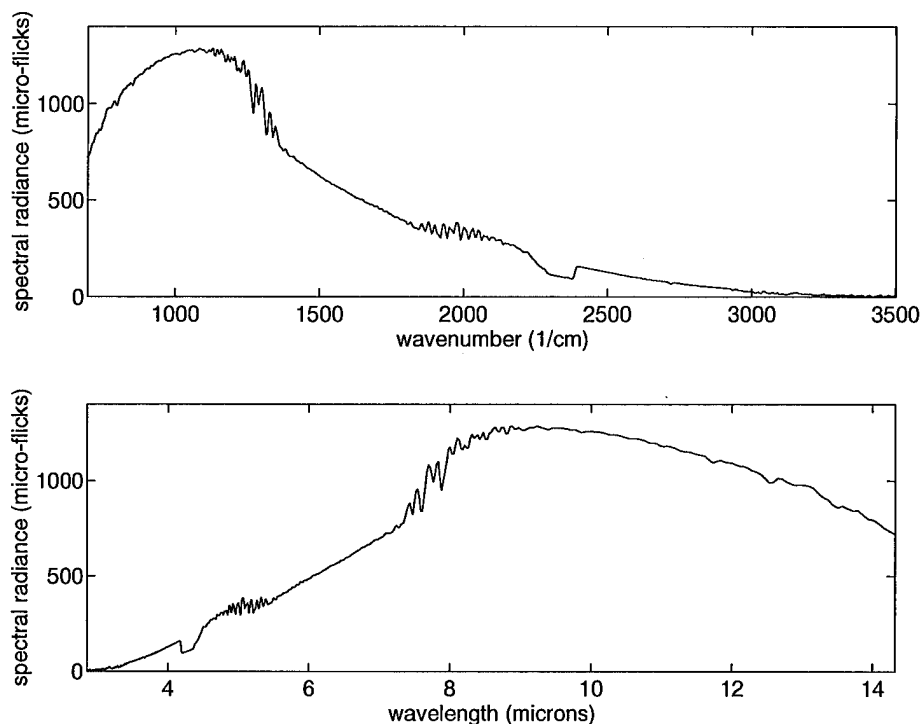
Figure 5. Typical spectrum measured by the Bomem FTS. This plot shows the same data plotted as a function of wavenumber (top) and wavelength (bottom). As collected, the 728 sample points in each spectrum are evenly spaced in wavenumber.

3.2.2 'Symptom Slew' Data Collection Configuration. The JMSP database represents a large and diverse collection of spectral signatures, including observations from four different geographic locations. While the detection software developed is able to process all this data, particular focus was given to one of the data collection sets. This set of observations, called the 'symptom slew' data set, consists of a set of 137 experiments conducted 6-10 June 1994 at Wright-Patterson AFB, Ohio. This data set contains over 2,700 spectral signatures of targets and over 1,500 non-target spectral signatures. All observations were made with the Bomem FTS operating from the Wright Laboratory Avionics Directorate testing tower in Area B of WPAFB. During this testing, the sensor field of view at the target was always smaller than the target dimensions.

	Ground Truth Name	Class	Target Description
1	Scud_B_Missile	target	SS-1 'Scud' missile
2	MAZ543_Launcher	target	launcher vehicle for Scud missile
3	Lance_Missile	target	U.S. MGM-52 short range missile
4	M752_Launcher	target	launcher vehicle for Lance missile
5	M35_Truck	target	U.S. military truck
6	M50_Tank	target	U.S. military artillery tank
7	Scud_B_Decoy	target	physical decoy for Scud B missile
8	MAZ543_Decoy	target	physical decoy for Scud missile launcher
9	CARC_Green_Panel	target	aluminum panel coated with U.S. Army CARC 383 green vehicle paint
10	UHaul_Truck	target	civilian utility vehicle
11	Tree_Canopy	non-target	trees surrounding test area

Table 1. Symptom slew data set target and non-target listing.

This data set serves as a good design and testing set since it contains a variety of targets along with a large set of non-target observations (mostly tree canopy). The targets and non-targets studied extensively from the symptom slew data set, along with their numerical designations, ground-truth file names, and descriptions are shown in table 1.

During the five days of the symptom slew data collection, the targets listed in table 1 were placed in various configurations to collect radiance measurements on targets located in the open, partially obscured, or in an operational state. The configurations of the targets listed in table 1 are important, as the radiance properties of the measurements made changed as the configurations were varied. The descriptions and the figures that follow offer a qualitative description of the testing configuration on each day of the data collection.

3.2.2.1 Day One: 6 June 1994.   The first day of data collection occurred on 6 June 1994. In order to baseline the measurements made in the following days, all targets were set in the observation area in front of the tree-line as indicated in figure 6. Observations began in the early morning hours (after midnight) and proceeded until mid-afternoon. All targets remained inactive during this testing

Figure 6.　Target configuration for symptom slew day one testing. Numbers in the diagram correspond to the target designations given in table 1.

(no engines running). Tree canopy observations were made by sampling the dense wooded area located behind the targets as indicated in figure 6

3.2.2.2　Day Two: 7 June 1994.　　Testing on day two of the data collection (7 June 94) supported the characterization of the changes in the radiance measurements when the targets were in an operational state (engines turned on). To accommodate these measurements, all targets remained in the same configuration as the previous day's testing (6 June 94), which is shown in figure 6. Tree canopy samples were collected from the same region of woods observed during the previous day's testing.

3.2.2.3　Day Three: 8 June 1994.　　On the third day of the data collection, the effects of target obscuration by the tree canopy were investigated. The basic target configuration was similar to that on 6 June 95 (see figure 6), except that the Scud and Lance missile assemblies (both missile and launcher) were placed within a small clearing in the surrounding woods such that both targets were partially obscured as viewed from the tower. In this location, the targets were in a radiant

Figure 7.   Target configuration for symptom slew day three testing. . Numbers in
the diagram correspond to the target designations given in table 1.

environment where the tree canopy contributed significantly to the measured target
spectral radiance, simulating a hidden target encounter. This configuration is shown
in figure 7. The engines on all targets remained off throughout the testing. Once
again, the tree canopy adjacent to the targets' location was sampled to provide
non-target observations.

3.2.2.4   *Days Four & Five: 9-10 June 1994.*   On the fourth and
fifth days of the data collection, the Scud assembly (missile and launcher), the Scud
decoy, the CARC (Chemical Agent Resistant Coating) green panel, and the Lance
assembly (missile and launcher) were moved to a distant hill that could be observed
from the tower, as shown in figure 8. None of the targets were obscured during this
testing, and all target engines remained off throughout the data collection. Non-
target observations were made by sampling the tree canopy near the targets on the
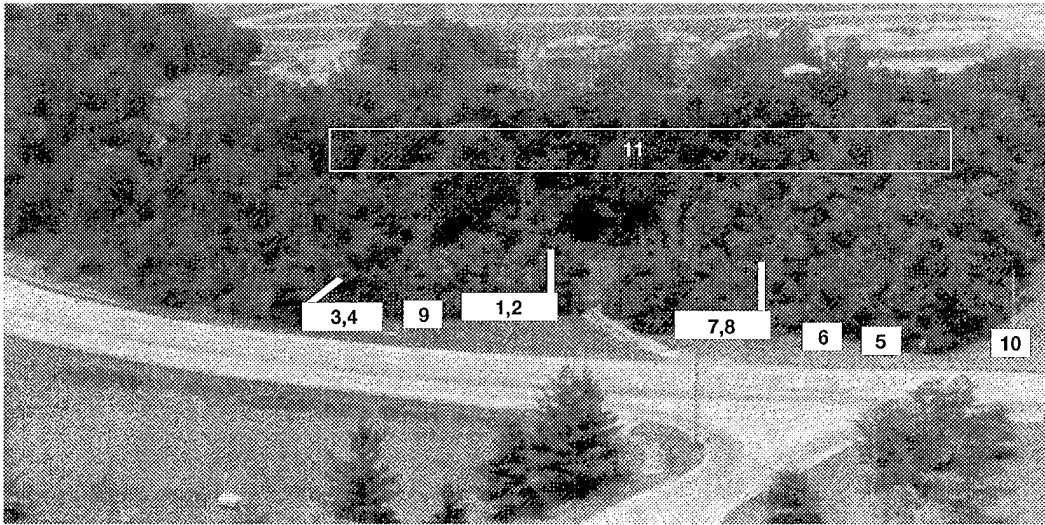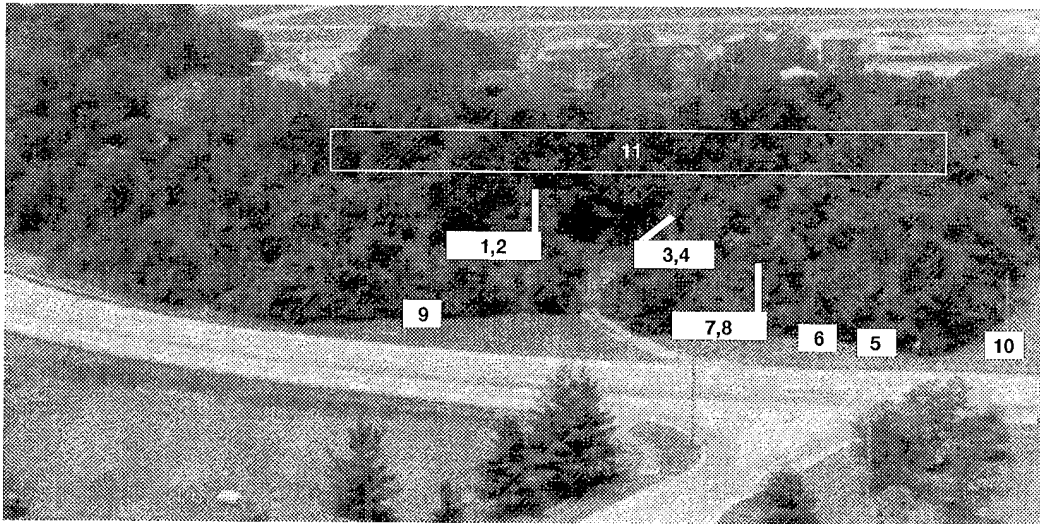distant hill.

Figure 8. Target configuration for symptom slew day four & five testing. Numbers in the diagram correspond to the target designations given in table 1.

3.3 Feature Set Development

As part of this research, a set of classification features was developed to facilitate target detection. The construction of these features was guided by review of the relevant literature as well as suggestions by the advisors to this research effort. A majority of the motivation behind the development of these features, however, came about through extensive analysis of the training data. Thus, the set of features developed reflects those features that were observed to offer adequate class separability within the limitations of the training set. In general, these features were not derived from first principles and then applied to the data, but rather, they were engineered using the relevant radiometric theory and the empirically observed statistical properties of the training data itself.

Within this section, the concept of atmospheric normalization, a technique developed during the course of this research, is introduced. In addition, each of the features designed, implemented, and incorporated into the detection software is discussed in detail. A definition of the feature is stated and a physical interpretation is offered where necessary.

37

3.3.1 Atmospheric Normalization. As was mentioned in the introduction to this section, a unique aspect of this research is the proposition and utilization of a concept referred to here as *atmospheric normalization*. In this technique, knowledge of the temperature of the atmosphere surrounding the target is used to calculate an atmospherically-normalized spectral radiance, $\hat{L}_\lambda$. Mathematically, this is defined as:

$$\hat{L}_\lambda(\lambda) = \frac{L_\lambda^{OBS}(\lambda, T_S)}{L_\lambda^{BB}(\lambda, T_A)}, \tag{14}$$

where $L_\lambda^{BB}(\lambda, T_A)$ is the blackbody spectral radiance as given in equation (5) evaluated at the temperature of the atmosphere T_A surrounding the observed target, whose temperature is T_S. As can be seen from equation (14), $\hat{L}_\lambda$ is a dimensionless quantity. Those familiar with radiometry should not confuse atmospherically-normalized spectral radiance with the normalized spectral distribution function f_λ (1:55), wherein spectral radiance is divided by the maximum spectral radiance reached by a blackbody at a given temperature. Instead, the atmospherically-normalized spectral radiance is similar to the spectral goniometric emissivity as given in equation (7), where the blackbody spectral radiance is calculated at the atmospheric temperature, T_A, as opposed to the observation temperature, T_S. Hereafter, atmospherically-normalized spectral radiance will be referred to as simply normalized spectral radiance.

As defined, computation of the normalized spectral radiance requires knowledge of the atmospheric or ambient temperature surrounding the target of interest. In practice, this information is rarely available in a direct form. In this work, information contained in the observed spectral measurement was utilized to estimate the ambient temperature.

As discussed in section 2.2.4, all infrared radiation is subject to absorption processes due to various molecular gases in the atmosphere. As is shown in figure 3 (see section 2.2.4), there is a large absorption band ranging from 5.5 μm to 7.3 μm due to vaporous H_2O. This absorption is near 100% around 6.25 μm (15). At this wavelength, therefore, radiation being emitted from the target of interest will

propagate only a few meters before being completely extinguished by the water vapor absorption. Thus, at 6.25 μm, the radiance measurements made do not reflect the spectral emission properties of the target of interest, but instead are a measure of the spectral emission properties of water vapor that is located a few meters from the sensor.

For a sensor located at an arbitrary distance from the target of interest, the spectral emission properties of water vapor a few meters from that sensor may not be indicative of the spectral emission properties of water vapor located around the target. For the training set used, in which the sensor was located several hundred meters from the target deployment site, the difference between the observed spectral emission and the actual spectral emission of the atmosphere surrounding the targets was small and therefore regarded as negligible. Thus, in order to estimate the ambient temperature at the target deployment site, the spectral radiance at spectrum sample number 232, which corresponds to a wavelength of 6.29 μm, was extracted from each spectral measurement. With this spectral radiance, the radiative temperature could be computed using equation (6). This radiant temperature was taken to be an estimate of T_A, the temperature of the atmosphere surrounding the target of interest, and was thus used to calculate the normalized spectral radiance as per equation (14).

The atmospheric normalization method described above is easily implemented given the spectral data available. The results that will be presented in Chapter IV demonstrate the significant performance improvement that may be obtained through its use. While this particular method has been proven beneficial, the practical implication of the concept of atmospheric normalization is that knowledge of the thermal background (by any means) is required to achieve high levels of target detection.

As will be seen in the feature definitions that follow, normalized spectral radiance is used extensively throughout the classification feature set. Ultimately, the utility of such features depends upon the ability to estimate the ambient tempera-

ture. Hence, the sensitivity of the classification performance to ambient temperature estimation errors is examined in detail in Chapter IV.

3.3.2 Feature Definition. All classification features developed in this work have been designed as *generalized features.* By 'generalized' feature, it is meant that the wavelength or wavelength band in which the feature is calculated is left as a parameter in the feature definition. By specifying values for these parameters, the generalized features become *specialized features* used for classification. Thus, these features may be applied to the MWIR as well as the LWIR. While the features may be applied in any region of the spectrum, there is no guarantee on the level of performance that may be achieved. For all of the generalized features, default wavelength or wavelength band values are defined. These defaults have been determined during the course of this research to yield satisfactory class separation under the Bayesian classification construct.

For the sake of clarity, throughout the feature definitions, either wavelengths (λ) or wavenumbers ($\bar{\nu} = \lambda^{-1}$) are used as the spectral designator. Remember that the spectral data is contained in an array in which the samples are evenly spaced in wavenumber. Whether the spectral designator is λ or $\bar{\nu}$, the corresponding spectral radiance sample does not change.

For the specialization of the spectral radiance, normalized spectral radiance, and the normalized radiant temperature features defined below, a default wavelength of 7.49 μm was chosen. This choice was made as it was observed during the feature development process that classification error rates were minimized at this wavelength for the normalized radiant temperature feature. Thus, 7.49 μm was selected as the default for the spectral radiance and normalized spectral radiance features as well. This wavelength selection allows for a direct comparison of the features' classification performance to that of normalized radiant temperature.

40

3.3.2.1 Spectral Radiance. The spectral radiance feature value, $L_\lambda(\lambda_S)$, is computed by extracting the spectral radiance at a particular wavelength λ_S directly from the observation spectrum. Any sample in the available spectrum may be used. The default is sample number 166, which corresponds to a spectrum wavelength of 7.49 μm. The feature value returned has the units of $\mu flicks$.

3.3.2.2 Normalized Spectral Radiance. The normalized spectral radiance feature, $\hat{L}_\lambda(\lambda_S)$, is calculated according to equation (14). As such, it represents the normalized version of the spectral radiance feature. The normalized spectral radiance feature values are dimensionless quantities. Feature values close to 1 indicate that the observed spectrum possesses radiance properties similar to a blackbody at the ambient temperature T_A at the spectrum sampling wavelength λ_S. The normalized spectral radiance feature may be calculated at any sample in the spectrum. The default sample for this feature is sample number 166, which corresponds to a wavelength of 7.49 μm.

3.3.2.3 Normalized Radiant Temperature. To compute the normalized radiant temperature feature value at a particular wavelength λ_S, the radiant temperature T_S is computed at the specified wavelength according to equation (6) using the observed spectral radiance. The ambient temperature T_A is then estimated by sampling the observation spectrum at 6.29 μm (λ_A) and using that spectral radiance measurement to compute the radiant temperature of the atmosphere. Using these numbers, the normalized radiant temperature $\hat{T}$ is then computed as:

$$\hat{T} = \frac{T_S}{T_A} \tag{15}$$

The default sampling wavelength for this feature is 7.49 μm (sample number 166). Normalized radiant temperature is a dimensionless quantity. As is the case with normalized spectral radiance, feature values near 1 indicate that the observation's

41

radiance properties at wavelength λ_S are similar to that of a blackbody at a temperature T_A.

3.3.2.4 *Normalized Band Radiance.*
The features defined thus far have made use of radiance measures at one particular wavelength (wavenumber). The normalized band radiance feature, however, uses spectral radiance measurements in a wavenumber band from $\bar{\nu}_1$ to $\bar{\nu}_2$. Within this band, the observed radiance of a source $L^{OBS}(T_S)$ at a temperature T_S may be computed by integrating the observed spectral radiance $L_{\bar{\nu}}^{OBS}(\bar{\nu}, T_S)$ with respect to $\bar{\nu}$ over the band:

$$L^{OBS}(T_S) = \int_{\bar{\nu}_1}^{\bar{\nu}_2} L_{\bar{\nu}}^{OBS}(\bar{\nu}, T_S) d\bar{\nu}.$$

Similarly, the radiance of a blackbody at the ambient temperature in the same band may be computed as:

$$L^{BB}(T_A) = \int_{\bar{\nu}_1}^{\bar{\nu}_2} L_{\bar{\nu}}^{BB}(\bar{\nu}, T_A) d\bar{\nu}.$$

Using these two band radiance measurements, the normalized band radiance is computed as:

$$\hat{L} = \frac{L^{OBS}(T_S)}{L^{BB}(T_A)} = \frac{\int_{\bar{\nu}_1}^{\bar{\nu}_2} L_{\bar{\nu}}^{OBS}(\bar{\nu}, T_S) d\bar{\nu}}{\int_{\bar{\nu}_1}^{\bar{\nu}_2} L_{\bar{\nu}}^{BB}(\bar{\nu}, T_A) d\bar{\nu}}. \tag{16}$$

The observations in the test set are expected to exhibit apparent temperatures near 300K, resulting in a spectral radiance peak in the 9-11 μm band (900-1100 wavenumbers). Hence, the default band for computing this feature was chosen to be 9.10-11.08 μm (samples 54-105). The integration over the band is accomplished using the trapezoidal rule (17:201). Normalized band radiance feature values are dimensionless quantities. Again, values for normalized band radiance that are close to 1 indicate that the radiance properties of the observation are similar to that of a blackbody in the band from $\bar{\nu}_1$ to $\bar{\nu}_2$.

3.3.2.5 Spectral Radiance Wavenumber Moment. Moments about the mean value $\bar{X}$ of a random variable X are called central moments. The n^{th} moment ($n = 0, 1, 2, \ldots$) is given the symbol μ_n^X and is calculated as follows (12:70):

$$\mu_n^X = E[(X - \bar{X})^n] = \int_{-\infty}^{\infty} (x - \bar{X})^n f_X(x) dx,$$

where $f_X(x)$ is a function which describes the manner in which X is distributed in the variable range (here, $-\infty$ to ∞). To describe how radiance is distributed over a band of wavenumbers for a spectral measurement, the spectral radiance at each sample in the band must be divided by the total band radiance. Hence, the spectral radiance distribution $f_{L_{\bar{\nu}}}(\bar{\nu})$ in a wavenumber band ranging from $\bar{\nu}_1$ to $\bar{\nu}_2$ may be computed as:

$$f_{L_{\bar{\nu}}}(\bar{\nu}) = \frac{L_{\bar{\nu}}(\bar{\nu})}{\int_{\bar{\nu}_1}^{\bar{\nu}_2} L_{\bar{\nu}}(\bar{\nu}) d\bar{\nu}}.$$

This spectral radiance distribution can be used to calculate the n^{th} wavenumber moment as follows:

$$\mu_n^{\bar{\nu}} = E[(\bar{\nu} - \langle \bar{\nu} \rangle)^n] = \int_{\bar{\nu}_1}^{\bar{\nu}_2} (\bar{\nu} - \langle \bar{\nu} \rangle)^n f_{L_{\bar{\nu}}}(\bar{\nu}) d\bar{\nu}. \qquad (17)$$

The expression given in equation (17) is used to compute the spectral radiance wavenumber moment feature value. In computing the distribution as well as the value of the moment, numerical integration is accomplished using the trapezoidal rule (17:201). To be consistent with the normalized band radiance feature, the default band that is used for calculating the wavenumber moment is between samples 54 and 105, which corresponds to a wavelength band from 9.10-11.08 μm. The default moment number for this feature is $n = 1$. Values of the spectral radiance wavenumber moment feature have dimensions of wavenumbers (cm^{-1}) to the n^{th} power.

3.3.2.6 Normalized Spectral Radiance Wavenumber Moment. Just as the spectral radiance distribution was used in computing the spectral radiance wavenumber moment, so too may a *normalized* spectral radiance distribution be computed in a band ranging from $\bar{\nu}_1$ to $\bar{\nu}_2$:

$$f_{\hat{L}_{\bar{\nu}}}(\bar{\nu}) = \frac{\hat{L}_{\bar{\nu}}(\bar{\nu})}{\int_{\bar{\nu}_1}^{\bar{\nu}_2} \hat{L}_{\bar{\nu}}(\bar{\nu})d\bar{\nu}}.$$

As before, this distribution may be used to define the n^{th} normalized spectral radiance wavenumber moment as:

$$\hat{\mu}_n^{\bar{\nu}} = E[(\bar{\nu} - \langle\bar{\nu}\rangle)^n] = \int_{\bar{\nu}_1}^{\bar{\nu}_2} (\bar{\nu} - \langle\bar{\nu}\rangle)^n f_{\hat{L}_{\bar{\nu}}} d\bar{\nu}. \tag{18}$$

When computing the value of this feature, all integrations are accomplished numerically using the trapezoidal rule (17:201). The default band used for calculating this wavenumber moment is from spectrum sample 54 to spectrum sample 105, corresponding to a wavelength band ranging from 9.10 μm to 11.08 μm. The default moment number for this feature is $n = 1$. These default values were chosen to be consistent with the spectral radiance wavenumber moment feature. Values of the normalized spectral radiance wavenumber moment feature have dimensions of wavenumbers (cm^{-1}) to the n^{th} power.

3.3.2.7 Spectral Radiance Correlation Distance. As was presented in section 2.2.3, for a given pair of wavelengths (λ_1, λ_2), the spectral radiance measurements of an ideal blackbody forms a straight line when the parametric curve $[L_\lambda^{BB}(\lambda_1, T), L_\lambda^{BB}(\lambda_2, T)]$ is plotted as a function of T. Thus, if a spectral observation possesses characteristics similar to a blackbody at the wavelength pair (λ_1, λ_2), then when plotted in the spectral radiance correlation plane, this observation should lie close to the blackbody correlation line. If the slope m and intercept b of the blackbody correlation line for a given wavelength pair is known, then the distance

44

from the point $(L_\lambda^{OBS}(\lambda_1), L_\lambda^{OBS}(\lambda_2))$ to that line is given by:

$$D_{(\lambda_1,\lambda_2)} = \frac{|L_\lambda^{OBS}(\lambda_2) - mL_\lambda^{OBS}(\lambda_1) - b|}{\sqrt{m^2 + 1}}. \tag{19}$$

The slope m and intercept b for the spectral radiance correlation line may be computed directly, once two points on the line are known. These points may be generated from the Planck radiation law, given in equation (5), and choosing two source temperatures, $T_1 = 280K$ and $T_2 = 330K$, which yield two widely-separated points in the correlation plane (see figure 2). These temperatures represent the range of apparent temperatures expected to be exhibited by the test data. Thus, with the blackbody correlation points, $(L_\lambda^{BB}(\lambda_1, 280K), L_\lambda^{BB}(\lambda_2, 280K))$, and $(L_\lambda^{BB}(\lambda_1, 330K), L_\lambda^{BB}(\lambda_2, 330K))$, the slope m and intercept b are given by:

$$m = \frac{L_\lambda^{BB}(\lambda_2, 330K) - L_\lambda^{BB}(\lambda_2, 280K)}{L_\lambda^{BB}(\lambda_1, 330K) - L_\lambda^{BB}(\lambda_1, 280K)}$$

$$b = L_\lambda^{BB}(\lambda_2, 280K) - mL_\lambda^{BB}(\lambda_1, 280K)$$

As can be seen from the distance formula in equation (19), the distance feature value is positive and is measured in $\mu flicks$. As a default, the band pair $(9.19\mu m, 11.03\mu m)$ is used in computing values for this feature. This band pair has been observed to yield high target detectabilities in the research efforts being conducted by the JMSP.

3.3.2.8 Normalized Spectral Radiance Correlation Distance. Instead of using the observed spectral radiance at two different wavelengths in the correlation distance feature, the normalized spectral radiance at those wavelengths can be used. When using the normalized spectral radiance as defined in equation (14), the task of computing the distance from an arbitrary observation point to the correlation line is simplified. The simplification comes from the fact that in the normalized spectral radiance correlation plane, the line of correlation is the line of slope $m = 1$ and

45

intercept $b = 0$ (the $y = x$ line). Thus, for the wavelength pair (λ_1, λ_2), the distance from the point $(\hat{L}_\lambda^{OBS}(\lambda_1), \hat{L}_\lambda^{OBS}(\lambda_2))$ to the correlation line is given by:

$$\hat{D}_{(\lambda_1, \lambda_2)} = \frac{|\hat{L}_\lambda^{OBS}(\lambda_2) - \hat{L}_\lambda^{OBS}(\lambda_1)|}{\sqrt{2}}. \tag{20}$$

The distance measure $\hat{D}_{(\lambda_1, \lambda_2)}$ is a dimensionless quantity, since normalized spectral radiance is also dimensionless. As is the case for its non-normalized counterpart, the default wavelength pair values are $9.19 \mu m$ and $11.03 \mu m$. Again, this band pair has been observed to produce high target detectabilities in research being conducted by the JMSP.

3.3.2.9 *Corrected Spectral Radiance.*

In the spectral signatures that are part of the training set, the predominant spectral characteristic of both target and non-target observations is emission which follows the general form of the Planck radiation law (a blackbody) with local spectral variations of differing relative magnitude caused by spectral emissivity differences in the target materials. To capture and exploit these local variations, it is necessary to remove the underlying blackbody spectral emission curve. As defined, the corrected spectral radiance feature accomplishes this task.

The motivation for computing this feature is provided by the definition of spectral emissivity, which is given in equation (7). In order to compute the spectral emissivity, one must know the temperature of the emitting source. An estimate of this temperature obtained by making radiant temperature measurements at several wavelengths and calculating the average of those measurements. Using equation (6), an estimate of the source temperature $\bar{T}_R$ can be obtained for N spectrum sampling points as:

$$\bar{T}_R = \frac{1}{N} \sum_{i=1}^{N} T_R(\lambda_i)$$

46

Using this estimate of the source temperature, the corrected spectral radiance L'_λ is defined as:

$$L'_\lambda(\lambda) = \frac{L_\lambda^{OBS}(\lambda, T_S)}{L_\lambda^{BB}(\lambda, \bar{T}_R)} \tag{21}$$

As implemented in the detection software, the estimate of source temperature is made by averaging over the radiant temperatures computed at each sample within a specified band. The corrected spectral radiance may then be evaluated at a particular wavelength (preferably within the band used to estimate source temperature). The default source temperature estimation band is from 8.00 μm (sample 144) to 12.00 μm (sample 36). This band was chosen for source temperature estimation as the radiant temperatures computed in the band are highly consistent. The default wavelength for corrected spectral radiance evaluation is 11.42 μm (sample 47), which was found to minimize the classification error rate associated with this feature.

3.3.2.10 Corrected Spectral Radiance Standard Deviation. Using the symptom slew training data, it was established empirically that the corrected spectral radiance of target observations exhibits larger variations with wavelength than do the corresponding non-target observations. As such, a classification feature was developed to quantify this variation. Over the spectrum range from 8.00 μm to 12.00 μm, it was observed that for both target and non-target spectra, the corrected spectral radiance maintained a value near 1. Furthermore, between 8.00 μm and 9.5 μm, target spectra exhibited much greater variation than in other ranges. Thus, in this range, the sample standard deviation of the corrected spectral radiance is computed as (18:46):

$$\sigma_{L'_\lambda} = \sqrt{\frac{1}{N-1} \sum_{i=1}^{N} (L'_{\lambda_i} - \bar{L}'_\lambda)^2} \tag{22}$$

The feature value computed is a dimensionless quantity (same as corrected spectral radiance). The wavelength band from 8.00 μm (sample 144) to 9.5 μm (sample 92) is the default range for computing this feature, taking advantage of the differences

between the target and non-target corrected spectral radiances observed in this band during the feature development process.

3.3.2.11 Energy Normalized Matched Filter. A linear space-invariant filter that is matched to a signal $s(x)$ has an impulse response given by $h(x) = s^*(-x)$ (8:178). Thus, if an input signal $g(x)$ is applied to a filter matched to $s(x)$, then the output $v(x)$ is given as:

$$v(x) = \iint_{-\infty}^{\infty} g(\xi)\, s^*(\xi - x)\, d\xi.$$

In application, the signal $s(x)$ is often called a template, while the filter is referred to as a *matched filter*. When constructing such a filter, the template $s(x)$ and the input signal $g(x)$ are often energy-normalized by dividing each sample of a signal by the square of the sum of the signal's components.

The output of such an energy normalized matched filter was employed as a classification feature in this research. In order to implement this feature, a template was constructed in a certain wavelength band by using the spectral radiance of a blackbody at the estimated temperature of the atmosphere surrounding the target of interest. The value of the feature for a particular observation is taken to be the maximum value of the matched filter output when the observation signal is used as an input. Additionally, the cross-correlation indicated in the definition of a matched filter is accomplished using a discrete Fourier transform method. Hence, the energy normalized matched filter feature value is computed as:

$$v_{\tilde{L}_\lambda} = \max[F^{-1}\{F\{\tilde{L}_\lambda^{OBS}(\lambda, T_S)\} \cdot F\{\tilde{L}_\lambda^{BB}(\lambda, T_A)\}^*\}], \tag{23}$$

where $\tilde{L}_\lambda^{OBS}(\lambda, T_S)$ is the energy-normalized spectral radiance of the observed target of interest and $\tilde{L}_\lambda^{BB}(\lambda, T_A)$ is the energy normalized blackbody spectral radiance at the ambient temperature T_A in the same band as the observed spectrum. For this

feature, the default wavelength band ranges from 9.10 μm (sample 105) to 11.08 μm (sample 54). This default band was chosen to be consistent with the other features in the set that utilize data in a particular band of the spectrum (e.g. normalized band radiance).

3.4 Class-Conditional Probability Distribution Function Generation

Essential to the Bayesian classifier is the estimation of the class-conditional probability distribution function for each class. In this research, discrete class-conditional probability distribution functions (PDFs) were generated for all features using a histogram approach with equally sized bins (3:424). The bins may be in a linear array, representing a single feature class-conditional PDF, or in a 2-D matrix, representing a two-feature joint class-conditional PDF.

The number of bins used for the PDF generation can be set arbitrarily by the user of the detection software. For the results reported here, however, default bin numbers built into the detection software were used; 100 bins for a one-dimensional PDF, 50 bins in each dimension for a joint PDF. For each bin, the probability density estimate was computed according to the relative frequency of occurrence, as stated in equation (11).

With a set of feature values X, the width of each bin is determined by the dividing the feature interval $[\min(X), \max(X)]$ into N_b bins of width w_b:

$$w_b = \frac{\max(X) - \min(X)}{N_b}.$$
(24)

A feature value x would be counted in the j^{th} bin if:

$$[\min(X) + (j-1)w_b] \le x < [\min(X) + jw_b].$$
(25)

For a joint (two-dimensional) distribution, a feature vector (x, y) would be counted in the (j^{th}, k^{th}) bin if:

$$[\min(X) + (j-1)w_{b_x}] \leq x < [\min(X) + jw_{b_x}] \ and$$
$$[\min(Y) + (k-1)w_{b_y}] \leq y < [\min(Y) + kw_{b_y}], \qquad (26)$$

where w_{b_x} and w_{b_y} are the bin widths in the x and y directions, respectively. With all bin frequencies counted, each bin in the i^{th} array or matrix was divided by the number of occurrences of class θ_i to obtain a relative frequency, which was taken as an estimate of the class-conditional probability within the range of feature values associated with each bin.

The feature interval $[\min(X), \max(X)]$ is determined automatically once the signature database has been parsed for the desired set of spectral observations and feature values for those observations have been computed. In order to obtain a more accurate estimate of the PDF without using more bins, the feature value interval can be 'cropped' by the user specifying any value for the feature minimum or maximum. If the minimum for cropping is less than the observed minimum, the observed minimum is used. Likewise, if the maximum specified for cropping is greater than the observed maximum, the observed maximum is used. While cropping can be used to increase the fidelity of the PDF estimate, it should be employed conservatively, as any observations that are mapped into a feature region outside the cropped feature interval will be classified according to their distance from the target or non-target distribution centroid (minimal distance decision rule), not according to the Bayes decision rule. Hence, a minimal average error is not guaranteed for those observations mapped to a feature region outside the specified interval.

Once a class-conditional PDF has been generated, the information is saved to a file to be used in performing detection runs on the test data. As constructed, the detection software allows a PDF file to be used for any set of test data as long as

the same classification feature set is employed. This allows distinct training sets and test sets to be employed in classifier testing.

3.5 Multiple Feature Classification

With a single feature, the method of quantifying classifier performance is clear: construct the class-conditional PDFs using training data, apply these distributions in a Bayesian classifier, and assess the results of that classification by comparing the results with known ground truth. If multiple feature values are computed for each observation, then the methods involved are not as clear.

Of particular concern is the issue of how the class-conditional probability may be quantified in a higher dimensional feature space. Instead of simply knowing the class-conditional probability for a particular feature x, $p_x(x|\theta_i)$, the joint class-conditional probability $p_X(x_1, x_2, \ldots, x_N|\theta_i)$ for a feature set X must be known. For N features in the set X, construction of N-dimensional class-conditional PDFs is required. In this research, construction of joint class-conditional PDFs above two dimensions was not attempted, as increasing the dimensionality of the feature space requires an increasing number of spectral observations in the training set in order to obtain a valid PDF estimate.

A feature space of higher dimensionality than two can be investigated, given that the features employed are statistically independent. In this case, the joint class-conditional probability can be determined by simply taking the product of the individual class-conditional probabilities. Thus, for N independent features x_j in the set X, $j = 1\ldots N$, the joint class-conditional probability distribution is given by:

$$p_X(x_1, x_2, \ldots, x_N|\theta_i) = \prod_{j=1}^{N} p_x(x_j|\theta_i). \tag{27}$$

In practice, the features developed are not likely to be completely independent, as they are ultimately based upon the same spectral radiance measurements.

Marginal performance improvement, however, may still be garnered by assuming that they are independent and computing the joint probability as indicated in equation (27). For each feature selected, a single-feature class-conditional PDF must be constructed. If the features chosen are not statistically independent, then sub-optimal performance will be realized from the Bayesian classifier, as the joint class-conditional probability has been determined incorrectly. While the performance may be sub-optimal, it may still be better than the performance realized using only a single classification feature.

3.6 Classification Testing Methods

As implemented in the detection software, the classifier may be tested on any set of data available in the JMSP database. The set of spectral observations to be tested is controlled by experiment and target lists, which are explained in detail in Appendix A. At the time that a test is initiated, a file which contains PDF information must also be supplied to the classifier. In general, the testing and training sets may be completely separate. The standard method used in this research, however, was to use the same set of observations for training and testing. In so doing, resubstitution and leave-one-out tesing are two well-established procedures by which the performance may be assessed.

When a classifier test is conducted, resubstitution methods are employed first. After performance has been assessed in this manner, leave-one-out testing is performed and a set of summary detection statistics is computed for comparison to the resubstitution results. When the leave-one-out testing is performed, the PDFs constructed when an observation is left out are computed similarly to the resubstitution PDFs. That is, the same number of bins, the same feature limits, and the same bin widths are employed. Thus, the resubstitution and leave-one-out results will form a set of consistent measurements that bound the Bayes error rate for the test set in the feature space being used.

As described in section 3.2.2, the symptom slew data set contains considerable variation in target configuration, which may also introduce variations in class separability. As a result, this set of experiments has been partitioned into two subsets: *unobscured* and *obscured*. The unobscured test set consists of observations made on all targets listed in table 1 during testing on days 1-2 and 4-5 of the symptom slew data collection. This set contains 793 target and 486 non-target observations for a combined total of 1,279 spectral signatures. The obscured test set consists of 297 observations made on the obscured MAZ543_TEL and M752_Launcher targets and 216 observations made on the surrounding tree canopy during day 3 of the symptom slew testing. Test results on the unobscured set should be considered as a less-stressful test of class separability whereas the obscured set is considerably more stressful. Neither set represents an extreme or rare target viewing condition.

3.7 Bayes Error Rate Estimation

As discussed in Chapter II, the Bayes error rate associated with a set of testing data in a given feature space may be estimated in several ways. With a given set of test data and a given feature, the estimate obtained will depend upon the method utilized in constructing the class-conditional PDF for each class. The Bayes error rate estimate is then the total error rate associated with tests performed using a particular PDF construction technique with a Bayesian classifier.

In this research, a discrete-bin density estimate (see section 2.3.3.1) is employed as the principle technique for PDF construction. The results presented in Chapter IV will reflect error rates established through classifier testing using the discrete-bin density estimate. In all cases, these error rates are substantiated with other non-parametric Bayes error rate estimates, specifically those obtained through k-nearest neighbor (k-NN) and Parzen window density estimates.

The method of error rate estimation using k-NN and Parzen density estimates is borrowed directly from previous AFIT thesis work in 1993 by Lieutenant Curtis

Martin (10). The MATLAB code developed by Martin (pknn.m) was incorporated without change into the software developed for this research. This pknn.m software takes in a set of feature value observations for each class in the test set (target vs. non-target) along with a range of k values for k-NN density estimation and a range of h values for the Parzen windows. The code then partitions the observation sets each into 10 independent sets and performs testing on each set for every value in the range of k (k-NN) or h (Parzen) specified. The results of these 10 tests are then averaged for each value of the k or h parameter and a set of Bayes error estimates for the k-NN and Parzen techniques is returned. The Bayes error for the observations in the test set is taken to be bounded by the minimal observed leave-one-out error estimate and the corresponding resubstitution result. The leave-one-out and resubstitution error estimates have confidence intervals associated with them, which will relax the bounding of the Bayes error rate. Hence, the Bayes error rate is bounded with 95% confidence between the lower limit of the resubstitution interval and the upper limit of the leave-one-out interval.

3.8 Standardized Performance Metrics

Performance levels achieved in classifier testing are quantified through the use of standardized metrics. With these metrics, performance levels may be objectively compared between tests conducted on distinct observation sets, or between tests performed on the same observation set using different classification features. These metrics are widely used in target detection work and have been employed successfully in many applications (13:94).

Once the classification testing is complete, the values for the following variables are computed for both the resubstitution and leave-one-out methods: the number of target observations, N_t; the number of target observations correctly classified, N_t^{corr}; the number of non-target observations, N_{nt}; and the number of non-target observations correctly classified, N_{nt}^{corr}. Using these variables, the following standardized

performance metrics are computed: The probability of target detection, P_d, given as:

$$P_d = \frac{N_t^{corr}}{N_t},\tag{28}$$

the false alarm rate per target declaration, FAR:

$$FAR = \frac{N_{nt} - N_{nt}^{corr}}{\left(N_{nt} - N_{nt}^{corr}\right) + N_t^{corr}}\tag{29}$$

and the total error rate or total probability of error, P_e:

$$P_e = \frac{\left(N_t - N_t^{corr}\right) + \left(N_{nt} - N_{nt}^{corr}\right)}{N_t + N_{nt}}\tag{30}$$

It is also possible to compute a 95% confidence interval on the total error P_e. In previous work (3:346), this interval was used to compare P_e for various detection algorithms. The 95% confidence interval about P_e represents the interval within which the total error rate will fall 95% of the time for a large amount of equivalently distributed data. To define this interval, the sample-based estimate of the total error rate variance must be compute as (3:347):

$$\sigma_e^2 = \frac{P_e(1 - P_e)}{N_t + N_{nt}}\tag{31}$$

It can then be shown that the 95% confidence interval about P_e is given by

$$[(P_e - 1.96\sigma_e), (P_e + 1.96\sigma_e)]\tag{32}$$

The 95% confidence interval on P_e was also computed for each detection run and included as part of the standardized performance metric set.

3.9 Summary

This chapter has offered a description of the methods employed in the target detection research to produce the results presented in Chapter IV. The topics discussed within the current chapter are summarized below:

- The spectral measurements from the JMSP database were discussed, including a description of the Bomem MB-100 FTS used to make field observations at Wright-Patterson AFB, Ohio on 6-10 June 1994. Characteristics of this data were presented, and the configuration for testing on each day of the collection effort was detailed.

- The concept of atmospheric normalization was introduced and justified. This concept is key to many of the classification features developed in this research. These features were defined and discussed in detail.

- The methods utilized for developing class-conditional PDFs through discrete-bin density estimation were presented. These PDFs may be used singly or in combination. When used in combination, the features are assumed independent, and the joint probability is taken to be the product of the individual feature probabilities. If independence is not assumed, the PDFs developed are two-dimensional, representing the true joint probability distributions for the two features employed.

- The standard methods for classification testing were outlined, including the manner in which resubstitution and leave-one-out tests were performed. Standard test sets for unobscured and obscured targets were discussed.

- The probability of target detection P_d, false alarm rate per target detection FAR, the total error rate P_e, and the 95% confidence interval were introduced as a set of standardized performance metrics. These metrics were used to quantify classifier performance and to compare that performance across classifier tests.

IV. Results and Analysis

4.1 Introduction

In the past chapters, the problem has been summarized, the relevant theory has been presented, and the methods for implementing that theory to solve the stated problem have been detailed. In this chapter, results obtained from classification testing using the symptom slew data are presented and analyzed. These results may be broken down into two types; performance measures and sensitivity analysis. When discussing performance measures, results will be presented in terms of:

- Single Feature Performance: Metrics detailing detection results obtained when a single feature is used in the classification process, indicating the degree of class separability that each feature offers.

- Multiple Feature Performance: Metrics detailing detection results obtained by assuming independence of features, and thus, computing the joint probability as the product of the individual probabilities.

- Joint Feature Performance: Metrics detailing detection results obtained by computing the true joint class-conditional probability distribution functions for various combinations of two classification features.

Using the classification feature that yields the top performance, the sensitivity of the Bayesian classifier employing that feature to various factors affecting detection performance will be investigated. Specifically, results will be presented for:

- Target Set Sensitivity: Variation in performance metrics as the testing target set is increased, indicating the relative performance of the classifier on each target tested.

- Target Configuration Sensitivity: Variation in performance metrics as the targets were placed in different configurations throughout the data collection.

- Diurnal Sensitivity: Variation in performance metrics as observation time changes.

- Mean Contrast Sensitivity: Variation in performance metrics as a function of the observed mean contrast metric (defined in section 4.9) in an experiment.

- Ambient Temperature Estimation Sensitivity: Variation of the total error rate as the estimate of ambient temperature becomes less accurate.

Taken as a whole, these performance results and analyses should provide an adequate picture of the ability to separate observations into target versus non-target classes using spectral-only infrared radiance data.

4.2 Single Feature Performance

In Chapter III, the feature set developed in this research effort was presented, with a description and mathematical definition being offered for each of the features. Of particular interest is the performance level that may be achieved on the testing data when each feature is employed individually. Quantifying these performance levels will indicate the degree of class separability that each feature provides when applied to the available spectral data. Investigation of the relative utility of employing these classification features concurrently will be reserved for a later section.

4.2.1 Feature Performance Ranking. Classifier tests were performed using the symptom slew unobscured standard test set. As described in Chapter III, this test set consists of spectral observations of all targets in the symptom slew target set taken between the hours of 8:00 a.m. and 6:00 p.m., on the first, second, fourth, and fifth days of the data collection effort, when no targets were obscured by the tree-line surrounding the test site.

For each of the classification features, class-conditional PDFs were constructed using the test set data. These PDFs were in turn employed in a Bayesian classifier, and the standardized performance metrics were computed using the results

of this testing. Feature performance was ranked in terms of probability of target detection, P_d (from highest to lowest). The ranking of the classification features is shown in table 2 and the corresponding performance metrics associated with each of these ranked features is listed in table 3 for resubstitution (R) and leave-one-out (L) testing.

As can be seen in table 3, the single classification feature that yielded maximum performance was the normalized spectral radiance wavenumber moment 1 calculated in the band between 9.10 and 11.08 μm, with a target detection rate between 94.70% (L) and 95.59% (R) and a false alarm rate between 3.56% (R) and 3.96% (L), corresponding to a total error rate of between 4.93% (R) and 5.71% (L). The feature that displayed the poorest performance was the spectral radiance correlation distance with the default waveband pair $(9.19\mu m, 11.03\mu m)$. Detection levels fall off significantly for features ranked between these extremes in performance. It is interesting to note that the false alarm rate is minimized using the normalized radiant temperature at 7.49 μm. While the detection rate for this feature is approximately 10% less than for the top-ranked feature, the corresponding false alarm rate, which is between 1.18% (R) and 2.90% (L), is almost half of that of the top-performing feature (normalized spectral radiance wavenumber moment 1).

The Bayes error rate for each of these single feature classification tests was estimated using the Parzen window technique and code mentioned in Chapter II and III (6, 10). In all cases except those indicated with an asterisk, the Bayes error estimate using the Parzen window technique for computing the class-conditional probability falls within the 95% confidence interval of the total error rate using the discrete bin density technique. The cases where the Parzen Bayes error estimate falls outside the interval may be considered suggestive of sub-optimal performance. However, since the 95% confidence interval on the Parzen Bayes error estimate is of the same magnitude as the interval given in table 3, significant overlap of the confidence intervals exists, indicating that the two error estimates are consistent for

Rank	Classification Feature Name
1	Normalized Spectral Radiance Wavenumber Moment 1: 9.10-11.08 μm
2	Corrected Spectral Radiance at 11.42 μm
3	Normalized Band Radiance: 9.10-11.08 μm
4	Normalized Spectral Radiance at 7.49 μm
5	Normalized Radiant Temperature at 7.49 μm
6	Corrected Spectral Radiance Standard Deviation: 8.00-9.53 μm
7	Normalized Spectral Radiance Correlation Distance: $(9.19\mu m, 11.03\mu m)$
8	Spectral Radiance Wavenumber Moment 1: 9.10-11.08 μm
9	Spectral Radiance at 7.49 μm
10	Energy Normalized Matched Filter: 9.10-11.08 μm
11	Spectral Radiance Correlation Distance: $(9.19\mu m, 11.03\mu m)$

Table 2. Single feature classification ranking based on P_d, highest to lowest.

Rank	Test	P_d	FAR	P_e	95% Confidence Interval	Bayes Error (Parzen)
1	R	0.9559	0.0356	0.0493	(0.0374,0.0611)	0.0305*
	L	0.9470	0.0396	0.0571	(0.0444,0.0698)	0.0516
2	R	0.9117	0.0334	0.0743	(0.0599,0.0886)	0.0774
	L	0.9117	0.0373	0.0766	(0.0620,0.0912)	0.0891
3	R	0.8815	0.0398	0.0962	(0.0800,0.1123)	0.0844
	L	0.8815	0.0398	0.0962	(0.0800,0.1123)	0.0906
4	R	0.8676	0.0269	0.0970	(0.0807,0.1132)	0.0734*
	L	0.8676	0.0471	0.1087	(0.0916,0.1257)	0.0804*
5	R	0.8449	0.0118	0.1024	(0.0858,0.1190)	0.0765*
	L	0.8449	0.0290	0.1118	(0.0945,0.1291)	0.0796*
6	R	0.8348	0.0741	0.1439	(0.1246,0.1631)	0.1188*
	L	0.8159	0.0849	0.1611	(0.1409,0.1812)	0.1282*
7	R	0.7516	0.0387	0.1728	(0.1521,0.1935)	0.1540
	L	0.7352	0.0395	0.1830	(0.1618,0.2041)	0.1626
8	R	0.7112	0.1350	0.2478	(0.2242,0.2715)	0.2424
	L	0.6759	0.1612	0.2815	(0.2568,0.3061)	0.2565*
9	R	0.6822	0.1426	0.2674	(0.2431,0.2917)	0.2799
	L	0.6419	0.1946	0.3182	(0.2927,0.3437)	0.2839*
10	R	0.6318	0.1180	0.2807	(0.2561,0.3053)	0.2737
	L	0.5864	0.2065	0.3511	(0.3249,0.3772)	0.2760*
11	R	0.6103	0.1597	0.3135	(0.2881,0.3390)	0.2798*
	L	0.5511	0.1739	0.3503	(0.3241,0.3764)	0.3283

Table 3. Performance metrics associated with single classification feature ranking.

each feature. The consistency of the error rates for all classification features demonstrates that the Bayesian classifier constructed has not been designed improperly. Thus, the computed detection, false alarm, and total error rates should be regarded as indicative of classification performance near the Bayes limit for this data and feature set.

4.2.2 Distribution functions of Top-Ranked Features. It is instructive to study the form of the class-conditional distributions for each of the top-ranked features in the classification feature list. A feature was considered as 'top-ranked' if its associated Bayes error rate was less than 10%. This criterion yielded a list of the top five features listed in table 2. When used in single feature classification, all of these features yield target detection rates greater than 84% with false alarm rates less than 5% (some yield significantly better performance).

The class-conditional PDFs for the normalized spectral radiance wavenumber moment 1 feature as computed in the 9.10-11.08 μm wavelength band is shown in figure 9. As is depicted in this figure, the target observations appear to be distributed in an approximately normal fashion with a mean value near 2 wavenumbers. On the other hand, the non-target distribution does not appear to be normal, but overlaps the target distribution minimally. In the feature space, non-targets tend to have higher moment values than targets. Additionally, the distribution of non-target observations is wider than that of targets. The probability of error indicated in figure 9 was computed by assuming that the test set would have equal a priori probabilities. Since this assumption does not hold for the test set at hand (60% target, 40% non-target), the actual probability of error is somewhat higher. This fact is true of all error probabilities displayed on the PDF figures in this section.

Figure 10 depicts the class-conditional PDFs measured for the corrected spectral radiance at 11.42 μm. As is seen from the figure, the non-target distribution is highly bi-modal, exhibiting two strong density peaks at feature values less than 1.
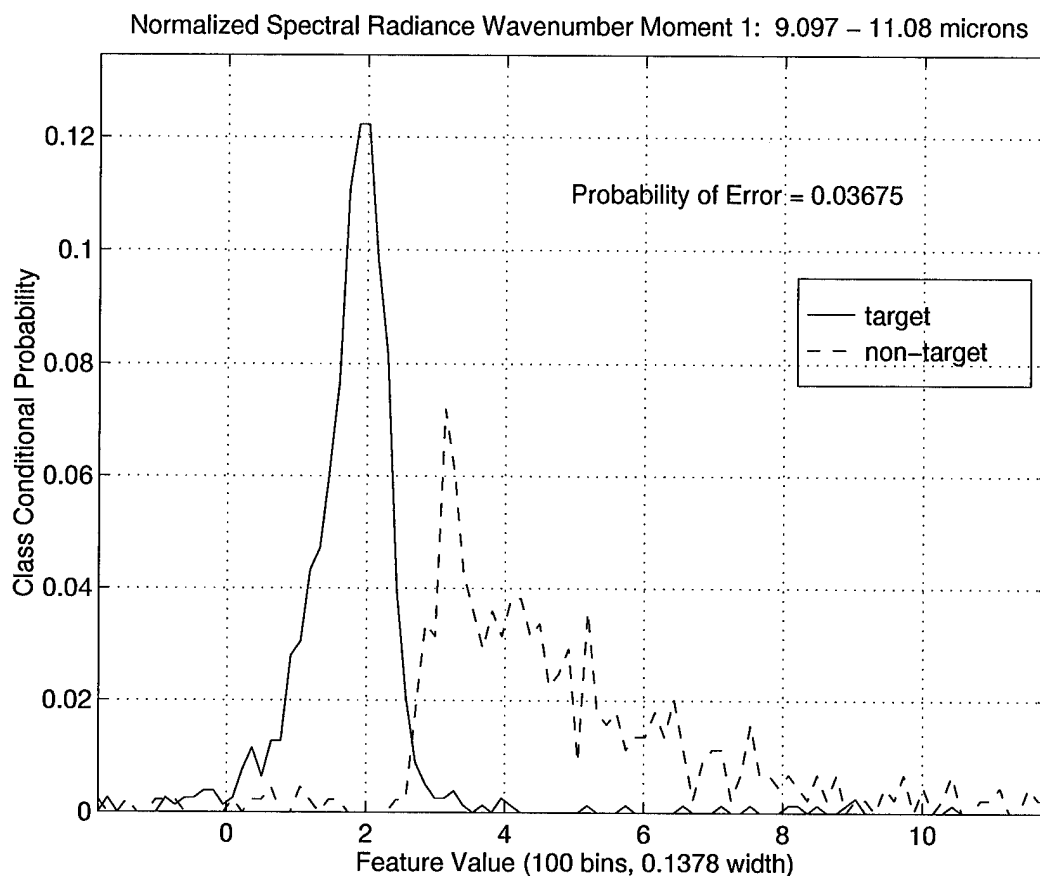
Figure 9. Class-conditional PDFs for Normalized Spectral Radiance Wavenumber Moment 1.

Nevertheless, as the bulk of the target distribution lies at feature values greater than 1, the amount of overlapping of the class-conditional distributions is sufficiently low to allow for single feature error rates that are the second lowest of all the features in the feature set.

The third-ranked classification feature was normalized band radiance in the 9.10-11.08 μm waveband. The class-conditional PDFs for this feature are shown in figure 11. In the space associated with this feature, the non-target observations appear to be highly confined, existing in a range of feature values (normalized band radiance) in the interval from around 1.0 to 1.1. Target observations tend to have higher values for the normalized band radiance. Additionally, the feature values

62
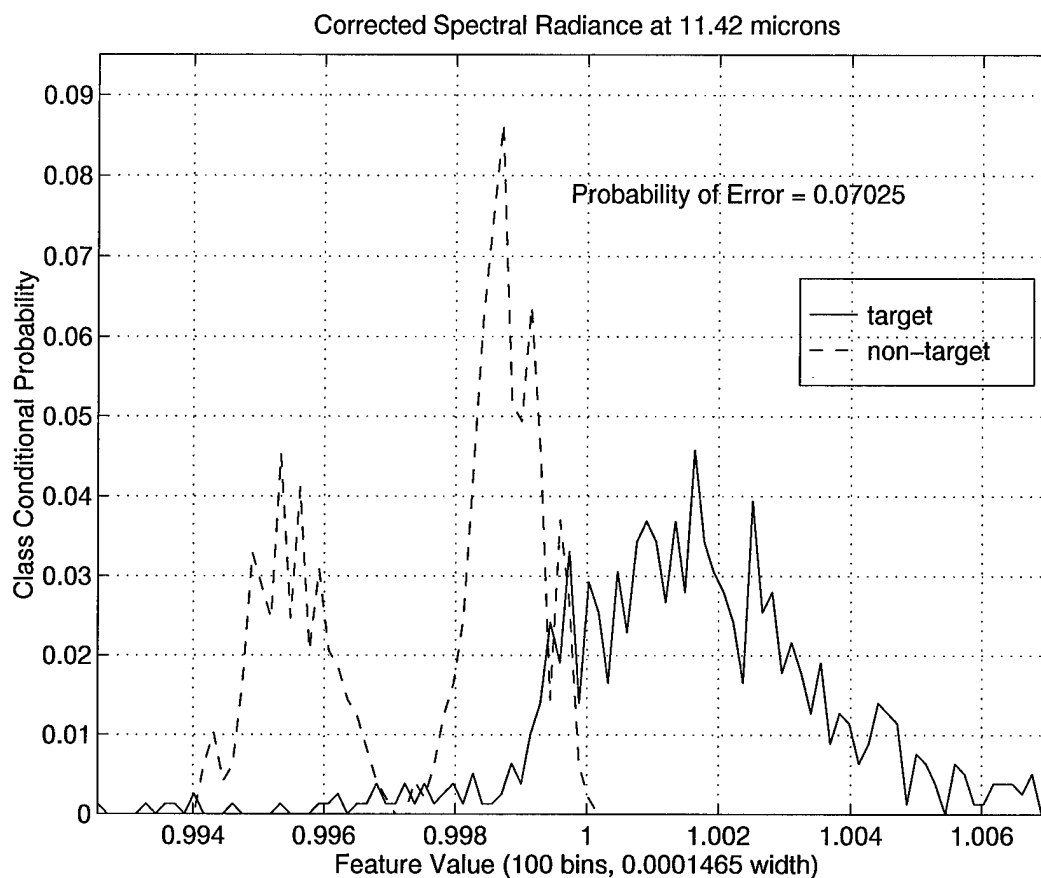
Figure 10. Class-conditional PDFs for Corrected Spectral Radiance.

associated with target observations have a wider range. Thus, the distributions overlap minimally, leading to a decreased error rate.

The normalized spectral radiance class-conditional PDFs, shown in figure 12, exhibit distributions similar to those associated with the normalized band radiance feature. Once again, the normalized spectral radiance values for non-targets tend to be lower (closer to 1) with less variance than the values for targets. Again, the confinement of the non-target distribution relative to that of the target observations provides for a low total error rate.

The normalized radiant temperature feature is computed in a manner similar to that of the normalized spectral radiance. In fact, these two features simply characterize the radiance properties of the observed source at a particular wavelength
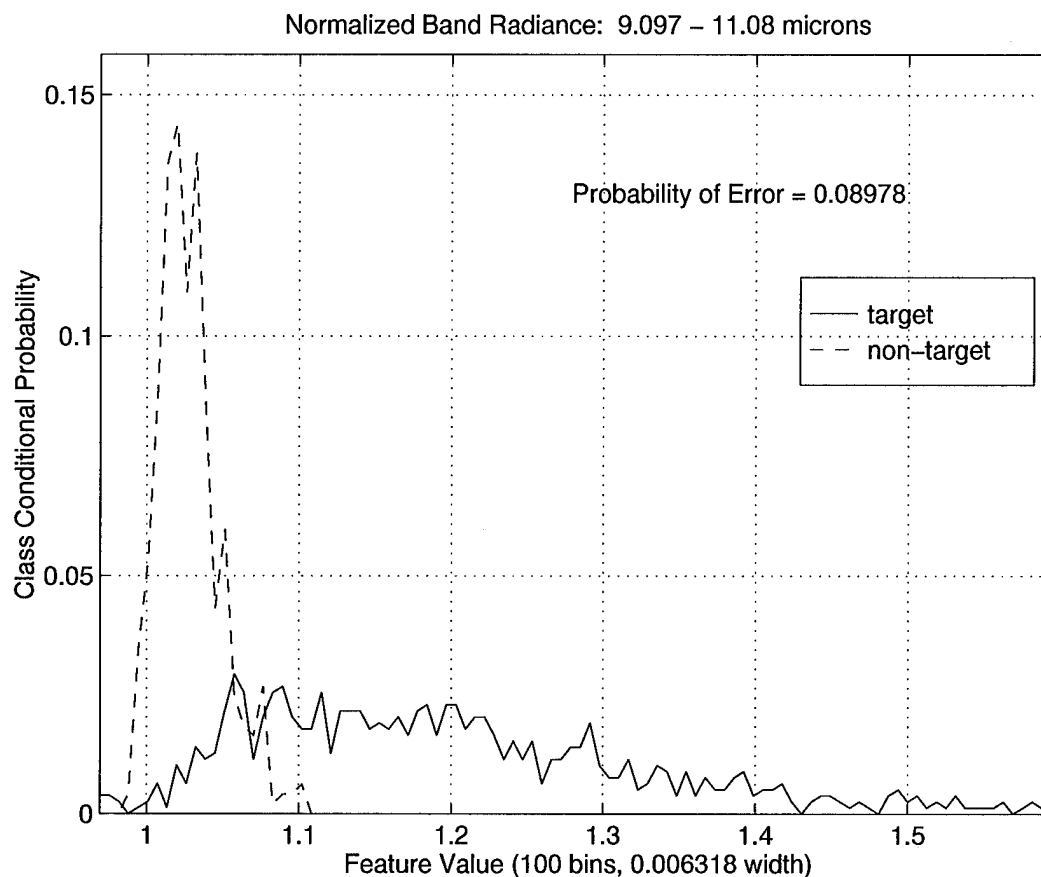
Figure 11. Class-conditional PDFs for Normalized Band Radiance.

in subtly different ways. Whereas the normalized spectral radiance is a ratio of the observed spectral radiance at a particular wavelength to the spectral radiance of a blackbody at the same wavelength but at the ambient temperature, the normalized radiant temperature feature instead uses a ratio of radiant temperatures. Thus, it is not surprising that the class-conditional PDFs for the normalized radiant temperature, seen in figure 13, display characteristics similar to those of normalized spectral radiance. Again, the feature values associated with non-targets have values close to 1, whereas the target feature values tend to be higher and more varied. This, of course, reduces the amount of overlap in the distributions, thus reducing the total error rate associated with classification processes which utilize this feature.
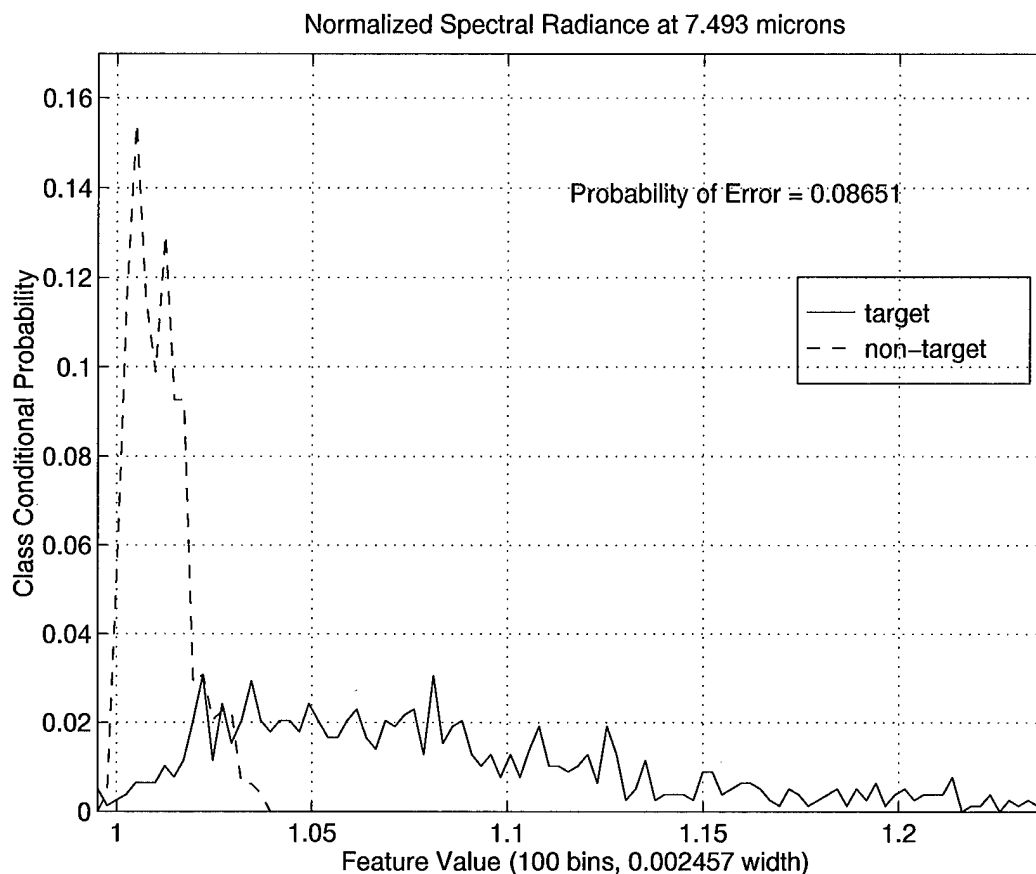
Figure 12. Class-conditional PDFs for Normalized Spectral Radiance.

4.2.3 Effect of Normalization on Feature Performance.   In the ranking of feature performance, it should be noted that four of the top five classification features are normalized features. The concept of atmospheric normalization, introduced in Chapter III represents one of the significant results of this research. Its utility comes from the significant increase in performance that is achieved through its application.

In the discussion of the class-conditional PDFs presented previously, the characteristic confinement of non-target versus target distributions is exhibited for many of the normalized features, indicating that non-targets display radiance properties that are similar to those of a blackbody at the ambient temperature surrounding them. In other words, these features seek to exploit the higher degree of equilibrium with the environment displayed by non-targets (especially vegetative). In fact, it has
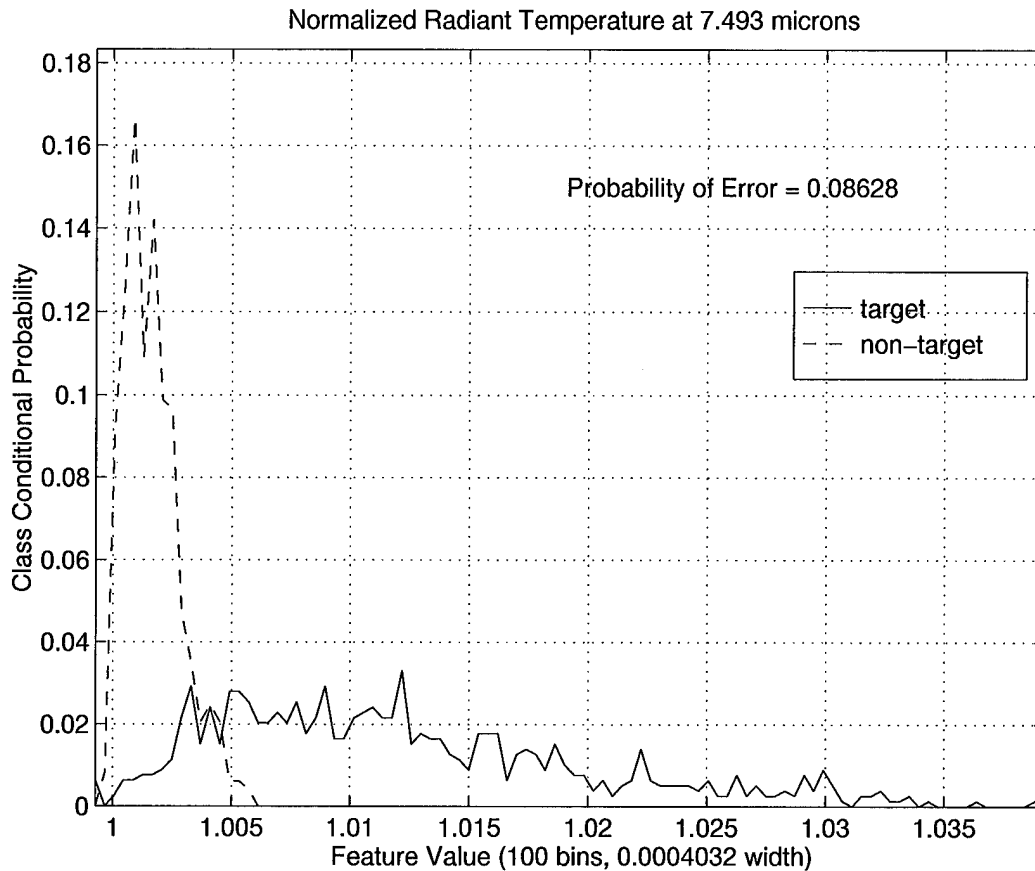
Figure 13. Class-conditional PDFs for Normalized Radiant Temperature.

been established that living vegetation employs active controls such as transpiration to maintain this equilibrium. Thus, it is reasonable that the radiance properties of non-targets are significantly less varied than targets once those properties are normalized using knowledge of the ambient temperature.

A convincing illustration of the utility of atmospheric normalization can be made by considering the class separation that is exhibited when normalized and non-normalized versions of the same type of feature are used. In the feature performance ranking, the ninth-ranked feature was that of spectral radiance at 7.49 μm. Its normalized counterpart, however, was ranked fourth. Furthermore, whereas the spectral radiance feature exhibited a Bayes error rate between 27.99% and 28.39%,

the normalized spectral radiance feature's Bayes error rate was bounded between 7.34% and 8.04%.

The origin of this large error disparity between the non-normalized and normalized versions of spectral radiance can be seen by comparing the class-conditional PDFs that are obtained when each is used as a classification feature. These PDFs are shown in figure 14. While pure spectral radiance measurements at 7.49 μm fail to confine or separate the two classes, simply normalizing the measurement to the expected spectral radiance for a blackbody at the estimated temperature of the atmosphere surrounding the source greatly increases the class separability. This is accomplished by first removing a large component of the scene-to-scene variation in the spectral radiance measurements. Once this variation has been removed, the ensemble distributions that remain reflect characteristics of the physical differences in radiance properties of targets versus non-targets. Hence, by atmospherically normalizing a feature, the separability of the classes is driven by characteristics which are directly related to actual class membership. In this light, therefore, it is not surprising that the process of feature normalization yields increased class separability.

4.3 Multiple Feature Performance

In section 3.5, the use of multiple features for classification was discussed. As was presented, the crux of the problem is to estimate the joint class-conditional probability for the two or more features being used. If the features are assumed to be statistically independent, then the joint probability can be obtained by taking the product of the class-conditional probabilities for each feature, as indicated in equation (27).

Acting upon the assumption of feature independence, combinations of the top-five ranked classification features in table 2 were used to classify the spectral observations in the symptom slew unobscured test set between the hours of 8:00 a.m.

67
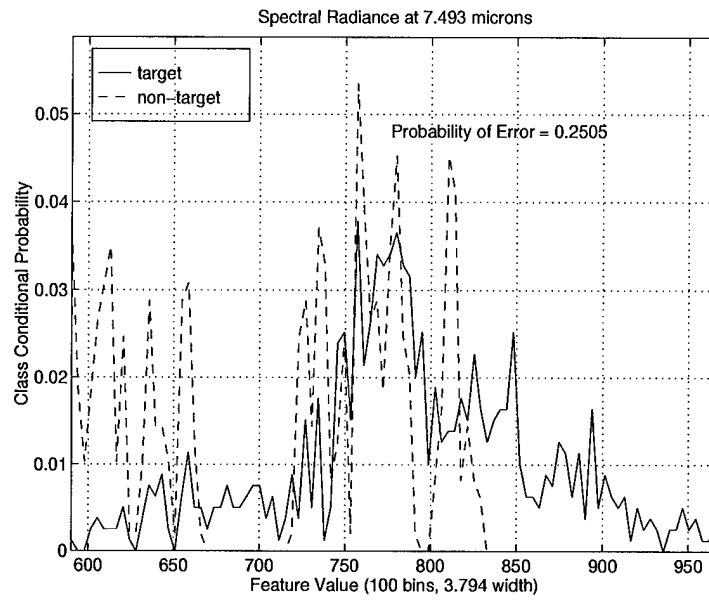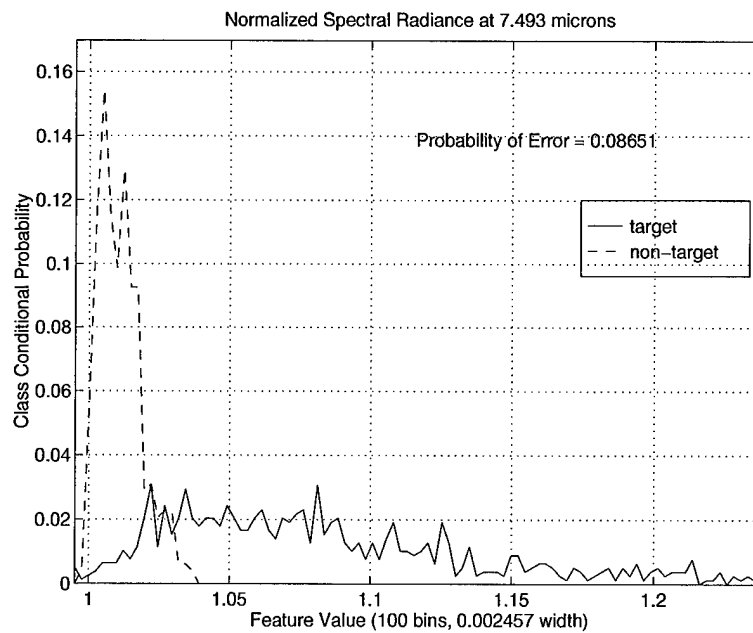
Figure 14. Effect of normalizing spectral radiance feature values. Comparison of the non-normalized PDFs (a) to normalized PDFs (b) for the spectral radiance feature.

Rank	Classification Feature Name
1	Normalized Spectral Radiance Wavenumber Moment 1: 9.10-11.08 μm
2	Corrected Spectral Radiance at 11.42 μm
3	Normalized Band Radiance: 9.10-11.08 μm
4	Normalized Spectral Radiance at 7.49 μm
5	Normalized Radiant Temperature at 7.49 μm

Rank	Test	Rank 1	Rank 2	Rank 3	Rank 4
2	R	0.0407	Rank		
	L	0.0508	2		
3	R	0.0336	0.0461	Rank	
	L	0.0477	0.0516	3	
4	R	0.0352	0.0446	0.0868	Rank
	L	0.0446	0.0516	0.0876	4
5	R	0.0336	0.0485	0.0923	0.0946
	L	0.0461	0.0555	0.1024	0.1095
Total Probability of Error, P_e					

Table 4. Total probability of error (P_e) using two-feature combinations, assuming feature independence.

and 6:00 p.m., as was done when using single features. The results of this testing are displayed in matrix form for re-substitution (R) and leave-one-out (L) tests in tables 4 and 5.

These results indicate that the feature combination that minimizes the total error rate under the assumption of feature independence is the top-ranked feature, normalized spectral radiance wavenumber moment 1, and the fifth-ranked feature, normalized radiant temperature. The total error rate for this combination of features is between 3.36% and 4.61%. The second and third top performing feature combinations are those using the normalized spectral radiance wavenumber moment 1 feature with the normalized band radiance and normalized spectral radiance features, respectively.

While the combination of normalized spectral radiance wavenumber moment 1 and normalized radiant temperature minimizes the total error rate, it does not

Rank	Classification Feature Name
1	Normalized Spectral Radiance Wavenumber Moment 1: 9.10-11.08 μm
2	Corrected Spectral Radiance at 11.42 μm
3	Normalized Band Radiance: 9.10-11.08 μm
4	Normalized Spectral Radiance at 7.49 μm
5	Normalized Radiant Temperature at 7.49 μm

Rank	Test	Rank 1	Rank 2	Rank 3	Rank 4
2	R	0.9672			
	L	0.9634			
3	R	0.9660	0.9344		
	L	0.9596	0.9319		
4	R	0.9660	0.9357	0.8827	
	L	0.9634	0.9306	0.8827	
5	R	0.9647	0.9294	0.8739	0.8689
	L	0.9622	0.9256	0.8663	0.8449
Probability of Target Detection, P_d					

Rank	Test	Rank 1	Rank 2	Rank 3	Rank 4
2	R	0.0328			
	L	0.0450			
3	R	0.0205	0.0094		
	L	0.0367	0.0160		
4	R	0.0230	0.0080	0.0251	
	L	0.0354	0.0147	0.0264	
5	R	0.0192	0.0081	0.0253	0.0241
	L	0.0366	0.0161	0.0351	0.0247
False Alarm Rate, FAR					

Table 5. Probability of target detection (P_d) and false alarm rate (FAR) using two-feature combinations, assuming feature independence.

concurrently maximize the probability of detection while minimizing the false alarm rate. The results in table 5 indicate that the probability of detection is maximized by the combination of the two top performing single features, normalized spectral radiance wavenumber moment 1 and corrected spectral radiance, with a detection rate between 96.34% and 96.72%. On the other hand, the false alarm rate is minimized by the combination of corrected spectral radiance and normalized spectral radiance, with a false alarm rate between 0.80% and 1.47%.

Comparing these results to those obtained through single feature classification does show an overall performance improvement when two features are used, even with the assumption of independence. Using combinations of two features that are assumed to be independent, the minimal total error rate dropped by more than 1% while the maximum probability of target detection increased by approximately 1% and the minimum false alarm rate observed dropped by close to 1%.

When working under the assumption of feature independence, no limit has been imposed on the number of features to be employed. Thus, the classification paradigm may also be applied equally well to combinations of three features. To effectively perform this analysis, the results obtained by pairing two features were used as a guide. As was discussed previously, the top three performing feature pairs utilized normalized spectral radiance wavenumber moment 1 paired with normalized radiant temperature, normalized band radiance, and normalized spectral radiance, respectively. Thus, to extend the classification paradigm to feature triplets, it is prudent to simply add another feature to these pairs. Specifically, the feature chosen to add to each pair was the feature that yielded the second greatest performance level when paired with normalized radiant temperature, normalized band radiance and normalized spectral radiance. Not surprisingly, the results in table 4 indicate that for all three of these features, the additional feature that produced this second greatest performance was corrected spectral radiance, the second-ranked single classification feature.

Rank	Classification Feature Name
1	Normalized Spectral Radiance Wavenumber Moment 1: 9.10-11.08 μm
2	Corrected Spectral Radiance at 11.42 μm
3	Normalized Band Radiance: 9.10-11.08 μm
4	Normalized Spectral Radiance at 7.49 μm
5	Normalized Radiant Temperature at 7.49 μm

Feature Ranking Combination	Test	P_d	FAR	P_e	95% Confidence Interval
1,5,2	R	0.9672	0.0217	0.0336	(0.0237,0.0435)
	L	0.9609	0.0379	0.0477	(0.0360,0.0594)
1,3,2	R	0.9672	0.0229	0.0344	(0.0244,0.0444)
	L	0.9634	0.0390	0.0469	(0.0353,0.0585)
1,4,2	R	0.9660	0.0254	0.0367	(0.0264,0.0471)
	L	0.9622	0.0390	0.0477	(0.0360,0.0594)

Table 6. Peformance obtained by combining corrected spectral radiance with top-three greatest performing feature pairs, with the assumption of feature independence.

The results obtained by pairing corrected spectral radiance with the top three feature pairs is shown in table 6. From these results, it can be seen that all feature triplets seem to yield approximately the same performance level in terms of total error rate, P_e, as witnessed by the large amount of overlap in the 95% confidence intervals. While the probability of target detection improves marginally ($< 0.25\%$) with three features, the false alarm rates appear to have risen slightly ($< 0.5\%$) and thus, the improvement of the total error rates is not statisically significant considering the associated 95% confidence intervals. Hence, it is not clear that utilizing more than two features in the paradigm of assumed independence produces any overall performance improvement.

4.4 Joint Feature Performance

It is apparent from the results presented in the previous section that by assuming feature independence, some performance improvement may be obtained. In

Rank	Classification Feature Name
1	Normalized Spectral Radiance Wavenumber Moment 1: 9.10-11.08 μm
2	Corrected Spectral Radiance at 11.42 μm
3	Normalized Band Radiance: 9.10-11.08 μm
4	Normalized Spectral Radiance at 7.49 μm
5	Normalized Radiant Temperature at 7.49 μm

Rank	Test	Rank 1	Rank 2	Rank 3	Rank 4
2	R	0.0360			
	L	0.0641			
3	R	0.0281	0.0399		
	L	0.0516	0.0821		
4	R	0.0274	0.0297	0.0915	
	L	0.0438	0.0711	0.1188	
5	R	0.0289	0.0360	0.0884	0.0876
	L	0.0477	0.0766	0.1063	0.0876
Total Probability of Error, P_e					

Table 7. Total probability of error (P_e) using two-feature combinations with joint class-conditional probability distribution functions.

the case of two feature combinations, the true joint probability may be calculated by constructing a two dimensional distribution function as described in section 3.4. Since no assumption is made regarding the degree of feature dependence, computing the true joint class-conditional probability should yield increased performance from the Bayesian classifier. The results of combinations of two features in classifying the symptom slew unobscured test set between the hours of 8:00 a.m. and 6:00 p.m. are shown in tables 7 and 8.

The results in table 7 show that the total decision error is minimized when the the top-ranked feature, normalized spectral radiance wavenumber moment 1, and the fourth-ranked feature, normalized spectral radiance are used jointly. The total error rate for this feature combination is shown to be between 2.74% and 4.38%. The second and third top performing feature combinations are those using the normalized spectral radiance wavenumber moment 1 feature with the normalized band

Rank	Classification Feature Name
1	Normalized Spectral Radiance Wavenumber Moment 1: 9.10-11.08 μm
2	Corrected Spectral Radiance at 11.42 μm
3	Normalized Band Radiance: 9.10-11.08 μm
4	Normalized Spectral Radiance at 7.49 μm
5	Normalized Radiant Temperature at 7.49 μm

Rank	Test	Rank 1	Rank 2	Rank 3	Rank 4
2	R	0.9685			
	L	0.9483			
3	R	0.9773	0.9470		
	L	0.9571	0.8815		
4	R	0.9760	0.9647	0.8701	
	L	0.9634	0.9054	0.8411	
5	R	0.9735	0.9483	0.8878	0.8878
	L	0.9584	0.9004	0.8613	0.8878
Probability of Target Detection, P_d					

Rank	Test	Rank 1	Rank 2	Rank 3	Rank 4
2	R	0.0266			
	L	0.0517			
3	R	0.0227	0.0118		
	L	0.0405	0.0155		
4	R	0.0203	0.0129	0.0199	
	L	0.0341	0.0218	0.0375	
5	R	0.0203	0.0066	0.0330	0.0316
	L	0.0355	0.0259	0.0367	0.0316
False Alarm Rate, FAR					

Table 8. Probability of target detection (P_d) and false alarm rate (FAR) using two-feature combinations with joint class-conditional probability distribution functions.

radiance and normalized radiant temperature features, respectively. It is interesting to note here that these same three feature combinations were also the top-performing combinations under the assumption of feature independence, though not ranked in the same order.

The same feature combination that minimized total decision error also maximized the probability of target detection when those features are used jointly. As is indicated in table 8, the probability of detection is maximized by the combination of normalized spectral radiance wavenumber moment 1 and normalized spectral radiance, with a detection rate between 96.34% and 97.60%. The false alarm rate, however, is minimized by the combination of corrected spectral radiance and normalized radiant temperature, with a false alarm rate between 0.66% and 2.59%.

In comparing the results obtained by employing two classification features jointly to those achieved using only a single classification feature, it is obvious that the increased dimensionality of the feature space provides for greater class separability. Specifically, the total error rate, P_e drops by nearly 2%, while the probability of target detection, P_d, increases by 2% and the minimal observed false alarm rate decreases by close to 1%. This overall increase in class separability can be understood more fully by considering the joint class-conditional PDFs for the top-ranked joint feature combination, as shown in figure 15.

As this figure illustrates, the two features utilized provide for increased class separability due to the fact that the distribution for each class is 'spread' in an orthogonal direction in the two-dimensional feature space defined by the feature combination. While the variation of feature values for targets is larger in the normalized spectral radiance direction than in the normalized spectral radiance wavenumber moment 1 direction, the opposite is true for non-targets. Thus, the area of overlap in the distributions is significantly decreased, leading to an increase in classification performance.
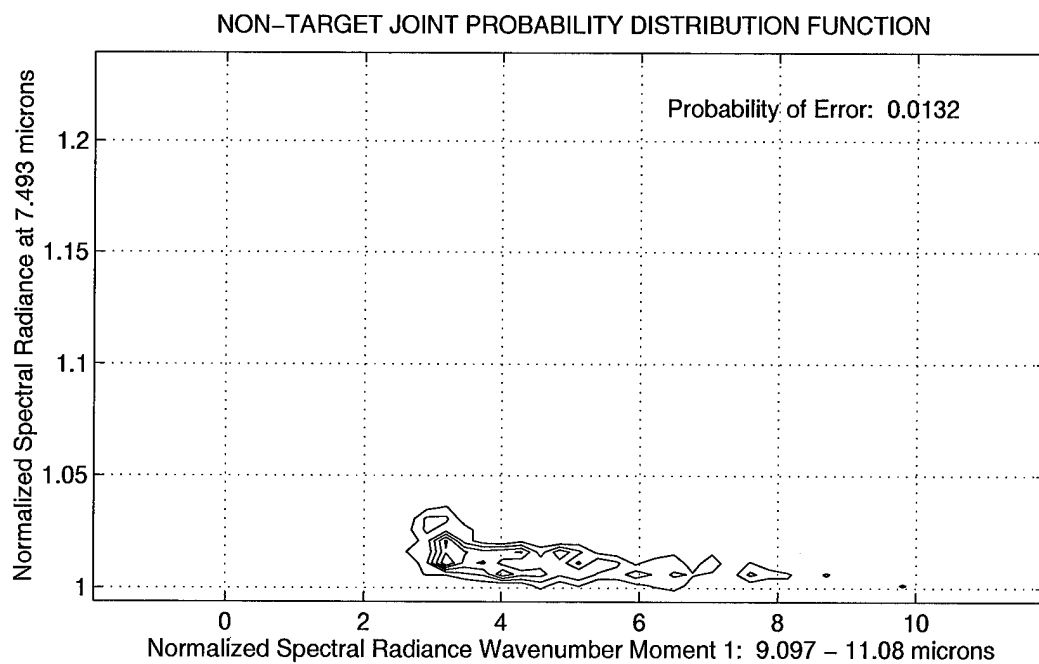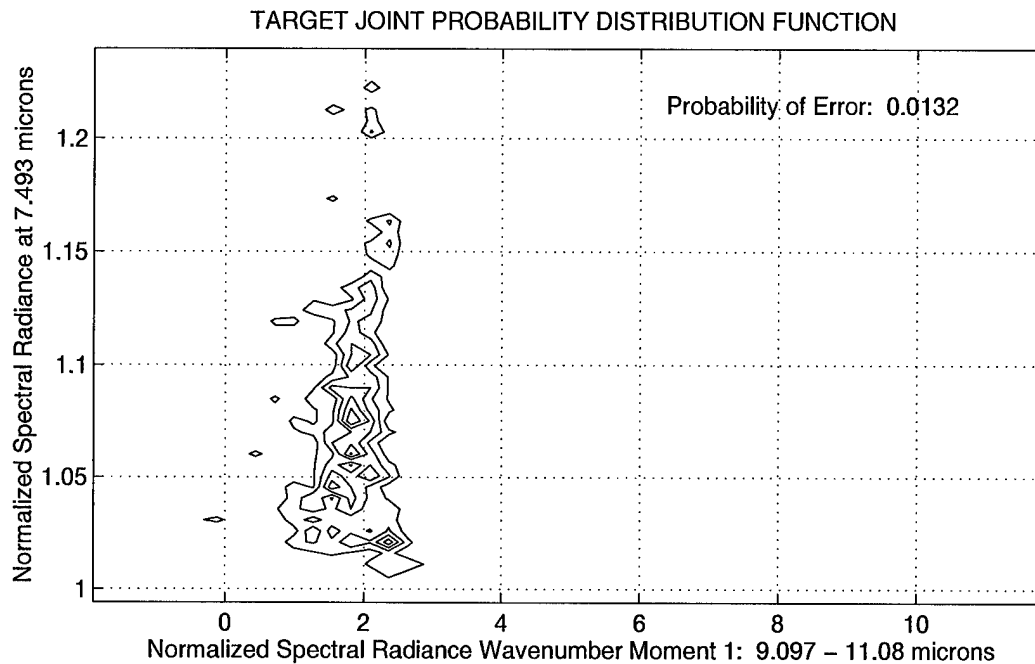
Figure 15. Joint class-conditional PDFs for the feature combination of normalized spectral radiance wavenumber moment 1 and normalized spectral radiance.

Classification Type	Test	Maximum P_d	Minimum FAR	Minimum P_e	95% Confidence Interval
Single Feature	R	0.9559	0.0118	0.0493	(0.0374,0.0611)
	L	0.9470	0.0290	0.0571	(0.0444,0.0698)
Two Features (assumed indep.)	R	0.9672	0.0080	0.0336	(0.0237,0.0435)
	L	0.9634	0.0147	0.0461	(0.0346,0.0576)
Three Features (assumed indep.)	R	0.9672	0.0217	0.0336	(0.0237,0.0435)
	L	0.9634	0.0379	0.0477	(0.0360,0.0594)
Two Features (used jointly)	R	0.9760	0.0066	0.0274	(0.0184,0.0363)
	L	0.9634	0.0259	0.0438	(0.0326,0.0550)

Table 9. Summary of best observed performance metrics for various classification methods.

4.5 Classification Performance Summary

The previous sections in this chapter have presented the performance results associated with testing on the symptom slew unobscured data set using features singly, in combination with the assumption of independence (both pairs and triplets), and jointly in a two-dimensional feature space. Each of these testing paradigms represent different methods for obtaining an estimate of the class-conditional probability, which lies at the heart of the Bayesian decision rule for class estimation. Perhaps the best way to summarize the results obtained and compare these methods is to look at the maximum P_d, minimum FAR, and minimum P_e with its associated confidence interval for each of the class-conditional probability estimation paradigms. These results are shown in table 9. From this table, it becomes obvious that moving from single-feature classification to two-feature classification is beneficial, with all performance metrics improving. It is less obvious, however, which method of employing two features is best. Although the resubstitution (R) values of the metrics for two features used jointly are significantly improved over those obtained under the assumption of feature independence, they vary more significantly under leave-one-out (L) testing. Ultimately, the high degree of overlap of the confidence intervals for these two methods does not make a compelling case for the joint use of classification

features over an assumption of independence. Marginal performance improvement for joint feature usage, however, is still implied by the results, indicating that the decision as to whether to use features jointly lies mainly in the relative difficulty of obtaining the joint class-conditional probability in the two-dimensional feature space.

When using multiple features under the assumption of independence, the results in table 9 do not indicate that using more than two features in the classification process will significantly improve performance. In fact, the results demonstrate that the use of three features may maintain the probability of target detection, but actually increase the false alarm rate. Since the 95% confidence intervals associated with two and three feature classification overlap, any improvements realized cannot be regarded as statistically significant. Thus, the use of more than two classification features under the assumption of independence is not demonstrated to be beneficial to the classification process.

Looking over the results obtained by single feature, multiple feature, and joint feature classification, one classification feature stands out as contributing significantly. The normalized spectral radiance wavenumber moment 1 feature (computed in the 9.10-11.08 μm waveband) was involved in minimizing total error while maximizing probability of target detection in all classification methods discussed. From these results, it appears to be a strong discriminant that can be effectively used in a non-imaging spectral infrared target detection system. Therefore, in the sections that follow, the sensitivity of this feature to several variables that would be encountered in such a system is explored and quantified.

4.6 Target Set Sensitivity

The results presented thus far are from testing the entire set of targets present during the symptom slew data collection. This list was presented in table 1, section 3.2. In table 3 (section 4.2), performance metrics were presented for resubstitution

Target Name	Spectra Classified	Correctly Classified	Incorrectly Classified	Percent Correct	Percent Incorrect
M35_Truck	96	96	0	100.0	0.0
UHaul_Truck	18	18	0	100.0	0.0
CARC_Green_Panel	11	11	0	100.0	0.0
Scud_B_Missile	8	8	0	100.0	0.0
Scud_B_Decoy	8	8	0	100.0	0.0
Lance_Missile	6	6	0	100.0	0.0
M50_Tank	101	100	1	99.0	1.0
M752_Launcher	63	62	1	98.4	1.6
MAZ543_TEL	388	365	23	94.1	5.9
MAZ543_Decoy	94	84	10	89.4	10.6
Tree_Canopy	486	458	28	94.2	5.8

Table 10. Breakout of classification performance on individual targets for re-substitution testing.

(R) and leave-one-out (L) testing using the normalized spectral radiance wavenumber moment 1 feature. While the results obtained reflect the level of classification performance using this feature, they do not indicate the level of performance achieved on each target individually.

Table 10 details the number of spectral observations for each target encountered in the symptom slew unobscured test set, along with the number of those observations classified correctly, incorrectly, and the corresponding percentages for resubstitution testing only. These results indicate that extremely high levels of performance may be achieved for a large set of the targets present in the test array. To understand the manner in which the performance metrics change as the target set is expanded, tests may be performed in which target observations are successively included in the test set, the class-conditional PDFs are recomputed, and classification is performed while the set of non-targets (tree canopy in this case) is kept constant. Such a set of tests produced the results detailed in table 11.

The results in table 11 show that the probability of target detection does not begin to fall off (from 1.0) appreciably until the MAZ543_TEL target is included in

Target Included	Test	P_d	FAR	P_e	95% Confidence Interval
M35_Truck	R	1.0000	0.2000	0.0412	(0.0251,0.0574)
	L	0.9896	0.2276	0.0498	(0.0322,0.0675)
UHaul_Truck	R	1.0000	0.1739	0.0400	(0.0243,0.0557)
	L	0.9912	0.1986	0.0483	(0.0312,0.0655)
CARC_Green_Panel	R	1.0000	0.1611	0.0393	(0.0239,0.0547)
	L	0.9920	0.1842	0.0475	(0.0306,0.0643)
Scud_B_Missile	R	1.0000	0.1529	0.0388	(0.0236,0.0540)
	L	0.9925	0.1750	0.0468	(0.0302,0.0635)
Scud_B_Decoy	R	1.0000	0.1455	0.0383	(0.0233,0.0533)
	L	0.9929	0.1667	0.0463	(0.0298,0.0627)
Lance_Missile	R	1.0000	0.1404	0.0379	(0.0230,0.0528)
	L	0.9932	0.1609	0.0458	(0.0295,0.0621)
M50_Tank	R	1.0000	0.1014	0.0381	(0.0243,0.0520)
	L	0.9879	0.1091	0.0450	(0.0300,0.0600)
M752_Launcher	R	0.9968	0.0828	0.0364	(0.0234,0.0494)
	L	0.9871	0.0890	0.0427	(0.0286,0.0567)
MAZ543_TEL	R	0.9614	0.0372	0.0447	(0.0330,0.0565)
	L	0.9571	0.0443	0.0515	(0.0389,0.0641)
MAZ543_Decoy	R	0.9559	0.0356	0.0493	(0.0374,0.0611)
	L	0.9470	0.0396	0.0571	(0.0444,0.0698)

Table 11. Variation in performance metrics as the target set is expanded to include the entire test set, using the normalized spectral radiance wavenumber moment 1 as a classification feature.

the target set. Furthermore, the false alarm rate continues to fall as the target set is expanded, since the inclusion of successive targets does not induce a significantly larger number of false alarms. This fact implies that the target distribution remains relatively constant as the target set is increased. While the normalized spectral radiance wavenumber moment 1 feature provides for significant class separability, it is not highly sensitive to target variations, making it an attractive feature to use for classification.

4.7   Target Configuration Sensitivity

Section 3.2 discussed the target configuration variations that are present in the symptom slew data set. Throughout the five days of data collection, the targets were placed in the open, obscured by the tree-line, moved further from the sensor, and observed in an operating state (engines turned on). These configuration variations are sure to introduce subsequent variations in the radiance data and hence, variations in the classification results when observations made in similar configurations are grouped together.

To test the target configuration sensitivity of the Bayesian classifier employing the normalized spectral radiance wavenumber moment 1 classification feature, the symptom slew experiments were placed into four different sub-groups. The experiments conducted on day one of testing, where all targets were placed in the open with their engines turned off, were taken as being one subgroup. Likewise, the observations taken on day two (targets in clear and engines running) and day three (some targets obscured by the tree-line) were each placed in a separate subgroup. Those spectra collected on day four and day five of testing (targets in the clear and distant from sensor) were placed together in the fourth subgroup.

The class-conditional PDFs associated with each of these subgroups were measured for the spectral data in the time interval between 8:00 a.m. and 6:00 p.m. These PDFs were then used with the Bayesian classifier and classification testing was conducted on each subgroup. Results from this testing are found in table 12.

The testing results from day one should be regarded as a baseline of classification performance, as all targets were located in the open in a non-operational state. As can be seen from the detection metrics, the probability of detection, having a value of 95.22% for resubstitution testing (R) and 92.49% in leave-one-out testing (L), is nearly equal to the probability of detection attained for testing on the entire unobscured observation set (see section 4.2.1). The probability of detection on day two improves significantly, as the targets observed during testing were placed in an

81

Observation Subgroup	Target Configuration	Test	P_d	FAR	P_e	95% Confidence Interval
Day 1	targets in clear engines off	R	0.9522	0.0822	0.0857	(0.0600,0.1114)
		L	0.9249	0.1424	0.1473	(0.1147,0.1798)
Day 2	targets in clear engines on	R	0.9873	0.0158	0.0196	(0.0069,0.0323)
		L	0.9873	0.0744	0.0632	(0.0409,0.0854)
Day 3	some targets obscured	R	0.6907	0.1079	0.2640	(0.2304,0.2977)
		L	0.6704	0.1490	0.3005	(0.2655,0.3355)
Day 4 and 5	targets in clear on distant hill	R	0.9730	0.0323	0.0301	(0.0126,0.0477)
		L	0.9405	0.0333	0.0466	(0.0250,0.0682)

Table 12.   Classification sensitivity to changes in target configuration using normalized spectral radiance wavenumber moment 1 as a single classification feature.

operational state, with the engines providing an internal source of heating for the target surfaces. Once again, on days four and five, the detection probabilities are similar to those obtained on day one.

The results obtained with the spectral observations collected during tests on day three, however, reflect significantly decreased classification performance. Recall that during these data collections, both the MAZ543_TEL and the M752_Launcher targets were placed within a wooded area at the target deployment site, thus being at least partially obscured by the surrounding tree canopy and providing a low-contrast radiant enviroment. As witnessed in the results of table 12, the probabilities of target detection dropped from around 95% to between 67.04% (L) and 69.07% (R).

The effects of obscuration on classification performance can be more fully understood by further dividing the subgroup of observations made on day three of the symptom slew data collection. Since the only targets obscured by the tree canopy on day three were the MAZ543_TEL and the M752_Launcher targets, the spectral observations associated with these targets were placed in a separate test set from the rest of the day three spectra. These sub-divisions of the data were classified independently and the results are presented in table 13.

Spectral Observations Tested	Test	P_d	FAR	P_e	95% Confidence Interval
Obscured	R	0.6700	0.2196	0.3002	(0.2605,0.3399)
	L	0.6094	0.2702	0.3567	(0.3153,0.3982)
Unobscured	R	0.9452	0.0676	0.0497	(0.0273,0.0721)
	L	0.9452	0.1375	0.0829	(0.0545,0.1113)

Table 13. Classification results obtained on obscured versus unobscured targets during day three of symptom slew data collection.

These results demonstrate that the performance achieved in testing the unobscured targets during day three are consistent with the performance for similar sets of data. For the obscured targets, however, the performance is significantly reduced, with low probabilities of detection and high error rates being realized. Thus, it is apparent that the effect of tree canopy obscuration on the radiant environment of the target can significantly reduce classification performance.

4.8 Diurnal Sensitivity

In the testing results reported thus far, the time interval between 8:00 a.m. and 6:00 p.m. was the standard testing interval from which spectral observations were calculated. It has been demonstrated that the normalized spectral radiance wavenumber moment 1 feature provides for a high level of class separability. At this point, it is an easy task to evaluate the class separability achieved using this feature at intervals throughout the entire period over which observations were made. This analysis is accomplished by dividing the total period of target observation into smaller time increments and constructing class-conditional PDFs within each of these intervals to be used in the Bayesian classifier. The spectral observations within each time interval may then be classified according to the standard methodology and the performance may be assessed through standard metrics.

In the symptom slew unobscured data set, the earliest observations of targets were made no earlier than 2:00 a.m. and no later than 4:00 p.m. Furthermore,

Time Period	Test	P_d	FAR	P_e	95% Confidence Interval
2:00-4:00 a.m.	R	0.7635	0.1484	0.2412	(0.2075,0.2748)
	L	0.7340	0.1902	0.2862	(0.2507,0.3217)
4:00-6:00 a.m.	R	0.7150	0.0738	0.2198	(0.1881,0.2516)
	L	0.6960	0.0844	0.2366	(0.2041,0.2692)
6:00-8:00 a.m.	R	0.7170	0.1128	0.2332	(0.2011,0.2652)
	L	0.6882	0.1534	0.2720	(0.2383,0.3058)
8:00-10:00 a.m.	R	0.9736	0.0537	0.0509	(0.0342,0.0676)
	L	0.9736	0.0837	0.0719	(0.0523,0.0914)
10:00 a.m.-12:00 p.m.	R	0.9771	0.0058	0.0177	(0.0068,0.0286)
	L	0.9685	0.0059	0.0230	(0.0106,0.0354)

Table 14. Variation of performance metrics during each two hour period between 2:00 a.m. and 12:00 p.m. for the symptom slew unobscured data set, using normalized spectral radiance wavenumber moment 1 as a classification feature.

only a small number of target observations were made after 12:00 p.m. (less than 20) and hence, any results obtained on such a small number of observations have little meaning. Thus, in order to test the time sensitivity of the normalized spectral radiance wavenumber moment 1 feature, the interval between 2:00 a.m. and 12:00 p.m. was broken down into five periods of two hours each. In each of these intervals, the observations were distributed in a fairly even fashion, such that in any two hour interval, approximately 350-400 target and 200-250 non-target observations were available for testing. Observation counts on this order provide for enhanced statistical significance of the subsequent distribution functions and the resultant classifier performance metrics. The performance metrics for each two hour interval are shown in table 14.

As in indicated in table 14, the probability of detection between the hours of 2:00 a.m. and 8:00 a.m. remains fairly constant between 70-75%. After 8:00 a.m., however, performance improves in a statisically significant fashion, with a probability of detection near 97% throughout the remaining hours of the morning. While only

84

limited data was available after 12:00 p.m., results from previous data collection analysis indicates that this level of performance would hold throughout the remainder of the daylight hours. These results imply that the success of the normalized spectral radiance wavenumber moment 1 feature depends upon the condition of either internal or external heating of the target surfaces such that those surfaces remain out of equilibrium with the surrounding atmosphere. Also implied by these results is that the surfaces need not be heated extensively before a high level of discrimination from natural objects is achievable. The data supports the conclusion that this feature would work well for detecting non-operating targets well after sundown and operating targets around the clock.

4.9 Mean Contrast Sensitivity

The previous two sections have dealt with the sensitivity of classification performance to target configuration and diurnal environment changes. Because there are many complex processes affecting the actual spectral emission of the targets under these varying conditions, the discussion presented has used mainly qualitative descriptions of the factors which actually lead to the sensitivities observed. It is possible, however, to quantify the effect of many of these factors in terms of a widely-used metric in infrared radiometry. This metric is commonly referred to as *mean contrast*. Often employed when dealing with infrared images, this metric seeks to express the 'observability' of a target as compared to its background. Mathematically, mean contrast, $\bar{C}$, may be defined as:

$$\bar{C} = \int_{\lambda_1}^{\lambda_2} \frac{|\langle L_\lambda^t \rangle - \langle L_\lambda^{nt} \rangle|}{\langle L_\lambda^t \rangle + \langle L_\lambda^{nt} \rangle} d\lambda, \tag{33}$$

where $\langle L_\lambda^t \rangle$ and $\langle L_\lambda^{nt} \rangle$ are the average spectral radiance associated with concurrent target and non-target observations, respectively. In this case, the averages are taken over a set of spectral observations, with the mean value of this set at each wavelength defining the average spectral radiance.
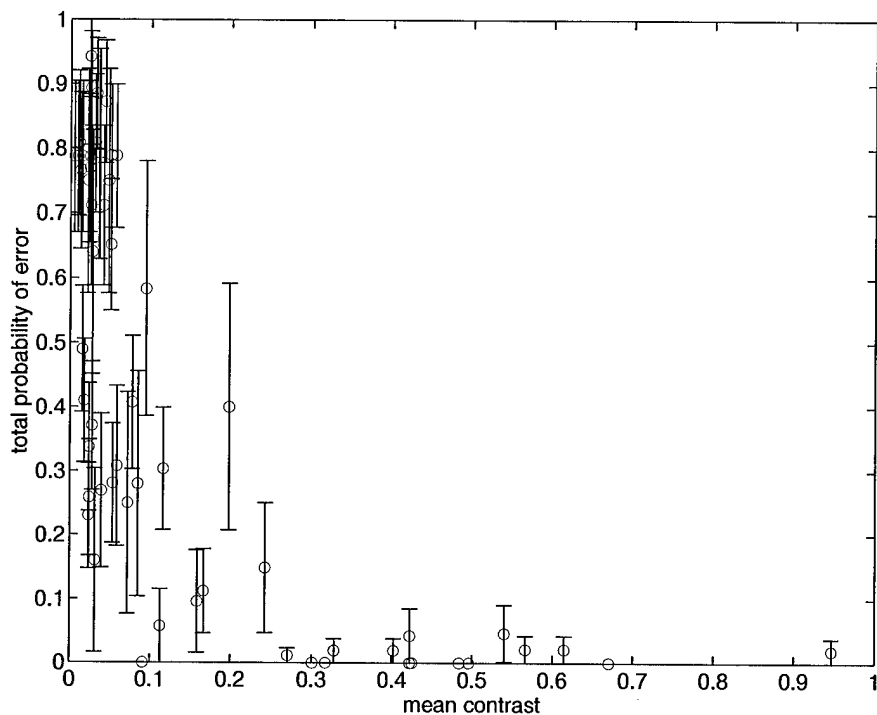
Figure 16. Sensitivity of total probability of error to observed mean contrast.

Using all experiments available in the symptom slew data set, tests of the classifier were conducted employing the normalized spectral radiance wavenumber moment 1 classification feature, with the mean contrast being computed for sets of observations associated with each experiment. Performance metrics for the classification processing on these observations were also computed. This data was used to produce the performance points shown in figure 16. In this figure, the error bars associated with each value of total probability of error represents the 95% confidence interval on that error rate.

The sensitivity results presented in figure 16 demonstrate that the observed target-to-background mean contrast behaves as a reasonable predictor of classification performance. The total error rate appears to vary with mean contrast in an intuitive sense, with increased performance levels being realized with increasing

contrast. As witnessed by this data, the error rate increases rapidly as the mean contrast drops below 0.1. It is also apparent from the these results that many of the experiments conducted in the symptom slew data collection effort can be accurately characterized as yielding 'low-contrast' spectral observations.

It is important to point out that while there is a definite relation between the total probability of error and the observed mean contrast, the classification performance cannot be completely parameterized in terms of mean contrast only. A parametric description of the classification performance achieved would have to include other variables as well. Thus, mean contrast should be used only as an *indicator* of the expected performance of the Bayesian classifier employed here.

4.10   Ambient Temperature Estimation Sensitivity

The effect of normalization on feature performance was presented in section 4.3. As the analysis showed, the atmospheric normalization technique employed by many of the classification features has a significant impact on their performance when compared to their non-normalized counterparts. In fact, in no case was a non-normalized feature ranked higher in classification performance than its normalized form (see table 2).

The success of this atmospheric normalization technique lies in the system's ability to estimate the temperature of the atmosphere surrounding the target. With the data available in the JMSP database, this could be done reliably by simply sampling the spectrum near a strong atmospheric absorption line. This method yielded a reasonable estimate of the ambient temperature due to the fact that the sensor was located only a few hundred meters from the targets. Hence, the atmosphere surrounding the sensor was at roughly the same temperature as the atmosphere surrounding the targets being observed. Keep in mind that due to atmospheric absorption, the emission observed near an atmospheric absorption line comes from the atmosphere within a few meters of the sensor.

87

In general, when operating from an airborne or space platform, the ambient temperature at the surface must be estimated. To explore the sensitivity of the normalized spectral radiance wavenumber moment feature to errors in the estimate of ambient temperature, deterministic errors were introduced to the ambient temperature estimation obtained through spectral sampling at 6.29 μm. The induced errors were quantified as a percentage of the estimated temperature (relative error = absolute error/ambient temperature estimate), and classification processing was performed with these errors present in the ambient temperature estimate. The resulting Bayes error rate (bound by the R and L curves) as a function of the relative error in the temperature estimate is displayed in figure 17 for relative errors from -3% to +3%, corresponding roughly to -10K to +10K absolute temperature error.

As is seen in figure 17, the error curve is not symmetric about the zero relative error point, with larger error rates being produced for large negative relative errors in the estimate than for the corresponding positive values. Furthermore, the minimum error is not precisely centered about zero, indicating that the method of ambient temperature estimation employed (spectral sampling) may itself be biased. Nevertheless, these results indicate that if the Bayes error rate is to be kept to less than 10%, then the relative ambient temperature estimate must be 1% or less. This roughly corresponds to an absolute tolerance within 3K of the actual ambient temperature.

4.11 Summary

This chapter has been used to present the performance results achieved by a Bayesian classifier employing the set of classification features developed as part of this research effort using spectral-only infrared radiance data. The sensitivity of the top performing feature to several factors that result in real-world data variability has also been investigated. These results have been analyzed and the principle discoveries are summarized in the list that follows:

PARZEN: Normalized Spectral Radiance Wavenumber Moment 1: 9.097 – 11.08 microns
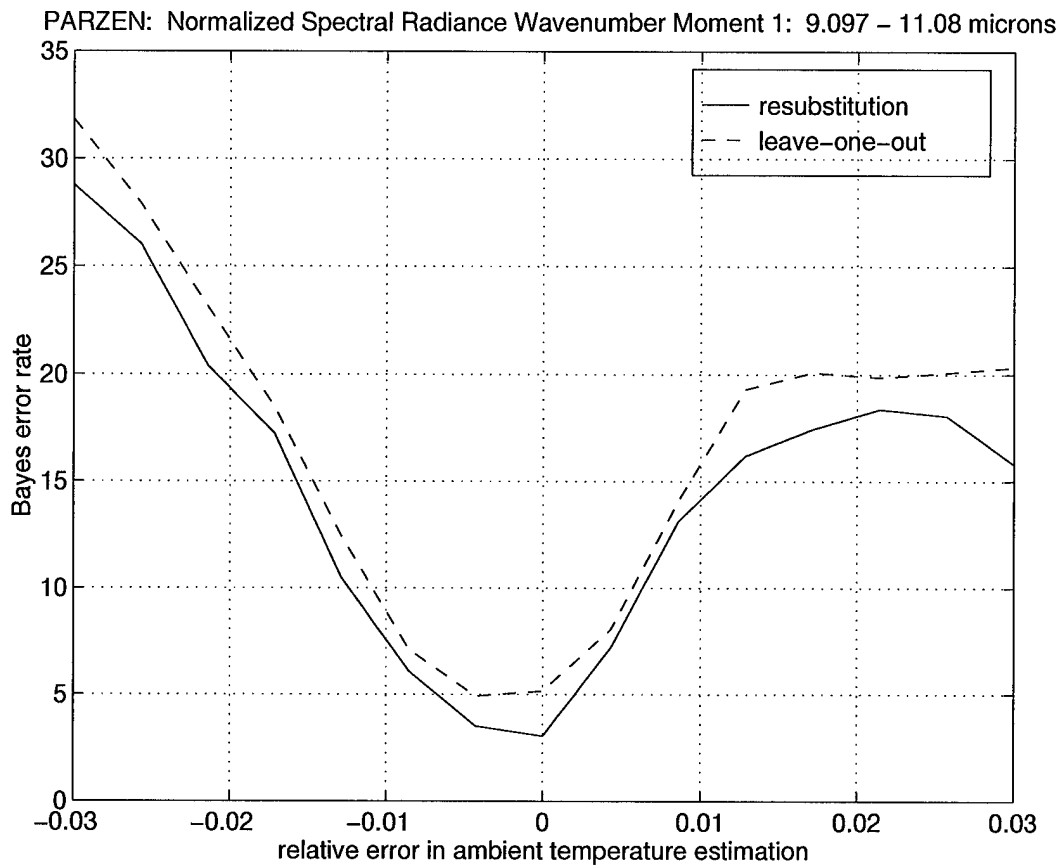
Figure 17. Ambient temperature estimation sensitivity of Bayes error rate for test-
ing on symptom slew unobscured data set for normalized spectral radi-
ance wavenumber moment 1 feature.

- Testing results obtained using each of the classification features separately
  indicate that the top-ranked single feature is the normalized spectral radiance
  wavenumber moment 1, which yields a total error rate of around 5% with a
  probability of target detection near 95%, and an associated false alarm rate of
  4%. Comparison of total error rates produced by the Bayesian classifier with a
  corresponding Bayes error estimate obtained through Parzen window density
  estimation demonstrates that the classifier developed here is operating in a
  near-optimal fashion.

- Atmospheric normalization of classification features significantly improves the
  performance of those features over their non-normalized counterparts. For

89

instance, the Bayes error rate associated with the spectral radiance at 7.49 μm drops from 28% to approximately 7.5% when this feature is normalized by estimating the ambient temperature surrounding the targets being observed.

- Multiple feature combinations were also employed, both with and without the assumption of feature independence. While combinations of two features enhance performance (increasing P_d by 1-2%), using more than two features was not demonstrated to be beneficial. In fact, when using three features under the assumption of feature independence, the false alarm rate increased while the probability of target detection remained nearly constant. In no case was statisically significant performance improvement obtained. Using two features jointly (measuring the joint PDFs) produced improved performance over the same feature combinations when assumed independent. Sensitivity of the PDFs to leave-one-out methods, however, result in a weaker bound of the Bayes error rate. Thus, the results obtained to not make a compelling case for measuring the joint PDFs over measuring the single feature PDFs and assuming feature independence.

- The sensitivity of classification performance to target set and target configuration variations was examined. These analyses reveal that most targets may be correctly classified at rates greater than 95% for all targets except the M752_Launcher and the MAZ543_TEL, which exhibited correct classification rates of 94% and 89%, respectively. When these targets were obscured by the tree canopy, the detection performance fell off significantly being reduced from over 90% to near 60%.

- Testing revealed that detection performance for the target set as a whole was reduced from near 95% after 8:00 a.m. to between 70-75% prior to 8:00 a.m. The detection response to external heating, however, was rapid, with detection rates greater than 95% observed in the time interval between 8:00-10:00 a.m.

- To quantify the radiometric observability of the targets in the test set, measurements of mean contrast were made for all experiments conducted in the symptom slew data collection effort. These measurements appear to be strong indicators of classification performance. The results of this sensitivity analysis reveal that error rates less than 10% require mean contrast measurements greater than 0.1.

- The sensitivity of classification performance to errors in the ambient temperature estimate used in feature normalization reveals that for the normalized spectral radiance wavenumber moment 1 feature, absolute temperature accuracies of approximately 3K are required to keep the total classification error rate below 10%.

These results, along with the theory and methodology presented in the previous chapters allows many conclusions to be drawn. These conclusions and recommendations for further research are discussed in the following chapter.

V. Conclusions

5.1 Introduction

The principle objective of this research was to quantify the probability of target detection, false alarms rate, and total error rate that may be achieved through specific automated classification processing of spectral-only infrared radiance data collected from time-critical mobile targets. This objective was motivated by the need to search wide areas in a timely fashion to discover the presence of these targets. Specifically, the research was intended to demonstrate that spectral-only infrared measurements could be used as a primary source of sensor data for target detection, and that the false alarm rate associated with this detection is sufficiently low such that the detection reports may be used as cues to another remote sensing system tasked with rejecting those false alarms. In this mode, a system utilizing this directed vision approach may be able to simultaneously produce high probabilities of detection ($P_d \geq 90\%$) and low false alarm rates ($FAR \leq 0.001$ per square kilometer).

In Chapter I of this thesis, the description of the problem at hand, a summary of the current knowledge, an outline of the scope, and the research approach were presented. Chapter II was used to introduce the relevant theory associated with the fields of radiometry and statistical classification. The methodology applied in this work was laid out in Chapter III, along with definitions of the classification feature set that was developed. Performance results and sensitivity analyses associated with the classification process were presented in Chapter IV. In the present chapter, the research accomplishments are summarized and conclusions are drawn from this work. Recommendations for future work in the area are also offered.

5.2 Summary of Developmental Activities and Research Advancements

The fulfillment of the objectives of this research required the development of software to search the JMSP database for a particular set of spectral observations, to

92

extract feature values from these spectra, compute class-conditional PDFs, perform Bayesian classification, estimate the Bayes error rate for a set of observations, and accomplish various types of sensitivity analysis. The resulting software system, which is described in Appendix A, meets all of these requirements and was used and refined extensively throughout the course of this work. Furthermore, the system provides an infrastructure upon which other research efforts may be built.

Along with the development of the software to implement the database searching, classification, and analysis tasks, a set of 11 generalized classification features were developed. These features are generalized in the sense that they may be computed in any waveband or at any wavelength and the performance of the resultant specialized features may then be assessed. This feature set uses radiance measurements at single wavelengths, multiple wavelengths, and integrated over certain wavebands. It also provides for alternative detection paradigms to the spectral correlation exploitation techniques currently being employed for spectral-only target detection.

During this research, the concept of atmospheric normalization of spectral radiance measurements was developed and implemented in creating many of the classification features in the established set. Atmospheric normalization uses knowledge of the ambient temperature of the atmosphere surrounding a target being observed in order to produce a feature value which maximizes classification performance by confining and separating target and non-target distributions.

5.3 Summary of Results Obtained

Classification tests conducted during the course of this research yielded a variety of performance results. Summarized here are the key results obtained using the symptom slew data set:

- Single feature classification results were used to rank-order the feature list, with the top-ranked single feature being the normalized spectral radiance wavenum-

ber moment 1, which yielded a total error rate of around 5% with a probability of target detection near 95%, and a false alarm rate of 4%. Parzen window density estimation techniques confirm that the Bayes error rate for this set of data using this feature lies between 3-5%.

- Atmospheric normalization of classification features improved the performance of those features significantly over their non-normalized counterparts.

- Multiple feature combinations may be employed for the classification process, both with and without the assumption of feature independence. Combinations of two features enhanced performance (increasing P_d by 1-2%), but using more than two features was not shown to be beneficial. Additionally, the results obtained do not make a compelling case for measuring the joint PDFs over assuming feature independence.

- Sensitivity analysis indicated that most targets present in the symptom slew data set may be correctly classified at rates near 95%. When these targets were obscured by tree canopy, the detection performance fell off significantly, being reduced to around 65%.

- Diurnal sensitivity analysis showed that a probability of detection of 70-75% may be achieved before thermal crossover. After crossover, however, detection rates increased rapidly and stayed at levels greater than 95% throughout the remainder of the daylight hours.

- The observed mean contrast was shown to be a reasonable indicator of classification performance. The results of sensitivity analysis revealed that error rates less than 10% require a target mean contrast greater than 0.1.

- While atmospheric normalization requires knowledge of the ambient temperature, some level of error is tolerable. Sensitivity analysis demonstrated that absolute temperature accuracies of approximately 3K are required to keep the

total classification error rate below 10% for the top-ranked classification feature.

5.4 Conclusions Drawn from Research

Spectral emissions from targets deployed in the field under the influence of a wide range of environmental and operational factors are dominated by thermal emission resembling blackbody radiation. Spectral emissivity variations, however, produce spectral features that, although sometimes subtle, may be exploited in order to provide class separability between targets and non-targets. In order to remove a large component of variability in environmental factors affecting the radiance properties of target and non-target objects, knowledge of the ambient temperature surrounding those objects must be garnered. This information may be used to normalize spectral radiance features, which has been shown in this research to increase class separability.

Classification testing reveals that for a set of approximately 1300 spectral measurements, target detection rates of 95-97% are readily achieved during daylight hours, with concurrent false alarm rates of 2.5-4%. Thus, for this set of measurements, the non-imaging spectral infrared target detection scheme has been demonstrated to meet the detection requirements for locating time-critical mobile targets. Furthermore, the concurrent false alarm rates achieved do not appear to limit the utility of such a sensor as the primary cuer in a directed vision system. Hence, an imaging system receiving cues from this sensor may be employed in a secondary mode to meet the overall false alarm rate requirements associated with time-critical mobile target location and prosecution.

5.5 Recommendations for Future Research

Many of the generalized features developed in this research effort were specialized by applying them at certain wavelengths or within certain wavebands that

produced reasonable class separation. This does not mean that the default specializations maximize the performance that may be achieved with these generalized features. Future research should therefore look toward discovering the wavelengths or wavebands at which these features produce a minimal total error rate. Furthermore, the set of classification features that exists now is by no means exhaustive. Thus, any future research conducted in this area would be well intentioned to implement and test new features to augment the current set.

The data utilized in producing the results discussed in this thesis were measured with a fully resolved target, meaning that the signal measured by the spectrometer was coming from target material only. If a sensing system is to utilize spectral signatures, there must be a trade-off between spectral and spatial resolution. Hence, it is a reasonable conclusion that targets may occupy a sub-pixel space, producing a spectral signature in which the target radiation is but a component. Thus, research conducted in the future should approach the issue of sub-pixel target detection, testing the adequacy of the established classification feature set for dealing with this type of data.

While the problem at hand was a two-class problem in which it was desired to classify all observations as being target or non-target, future research efforts should be aimed at increasing the number of classes in the problem. In so doing, the process may be regarded as moving from automatic target detection into the realm of automatic target recognition, wherein the declaration of tank versus truck, for instance, may be the goal.

From the results achieved in the current research, it is clear that spectral signatures offer a great deal of information that is not available in spatially resolved images. Furthermore, this information provides evidence of the physical radiative properties of the object being observed. This information should be regarded as complementary to the geometric radiance properties of the object that can be obtained from spatially-resolved image data. The two types of data can be made to work in

concert to indicate the class membership of the object being sensed remotely. With the conclusions of this research in hand, an important step toward this goal has been taken.

Appendix A. Infrared Spectral Target Detector Software

A.1 Introduction

Much of the developmental work in this research effort centered around the design and implementation of a classification software system to perform all tasks associated with producing the results reported in this thesis. This software remains a valuable tool for use in future work. Thus, it is important to outline here its functionality and operation. In this appendix, the main functions of the software are overviewed and key points involved in working with the classification feature set are presented. In addition, the classification processing output products are discussed.

A.2 Main Menu Functions

To enter the detection software environment, the user must simply type 'detect' at the MATLAB prompt. Upon entering the software, a working directory must be specified or the default working directory accepted. Once this has been done, the system main menu will appear offering many functions. The functions, which can be invoked by pressing their corresponding button on the main menu, are described in the sections that follow.

A.2.1 Specify Working Directory. This function permits the user to change the working directory for the session. The working directory is the directory from which files are read by the system as well as to where files are written. The system assumes that this directory is located within the directory from which MATLAB was initially called. If no directory name is specified, then the system uses a directory called 'default_working.'

A.2.2 Specify Printer Selection. With this function the user may specify to which printer hard-copy output is to be directed. The printer name, exactly as entered by the user, will be used in a MATLAB print -P command. If so desired,

printing output may be directed to a file. To do this, the word 'file' must be entered as the printer name. Once this specification is made, the system automatically prompts the user for an output file name when a print command is issued.

A.2.3 Compute Class-Conditional PDFs.   Selecting this option will allow the user to construct class-conditional PDFs from spectral observations in the JMSP database. In doing this, the user must first specify the dimensionality of the PDFs (1-D or 2-D) and then specify which features from the feature list are to be used. Upon completion of the database search, after the user has entered the number of bins to be used in each dimension of the PDF, the target and non-target distributions are displayed. These distributions may be printed to hard-copy if so desired. To make more efficient use of the histogram bins, the feature space may be cropped by specifying limits in each feature dimension. After the PDFs have been constructed, these measurements may be saved in a file for later use in classification processing.

A.2.4 Run Detection Algorithm.   Once this option is chosen, many parameters controlling the classification processing must be specified. Specifically, the user must select what type of classification will be carried out (single feature, multiple feature, etc.). After selecting the classification type, the user specifies the features to be employed along with the name of their corresponding PDF file (saved earlier upon constructing the PDF). Once the limitations of the database search are established, classification processing proceeds through the experiment list until all occurrences of the targets in the target list have been classified. At this point, the system automatically performs leave-one-out testing. The results are written to a file whose name is specified by the user.

A.2.5 Non-Parametric Bayes Error Estimation.   Choosing this option from the main menu allows the user to bound the Bayes error rate for a particular set of data and classification features using k-nearest neighbor (k-NN) and Parzen window

99

techniques. The user must therefore specify the limits in the value of k to be used in the k-NN technique as well as the range of values to be used for h in the Parzen technique. The number of points in this range of parameters must also be entered. Direct use of the Parzen/k-NN code developed by Martin (10) is made in producing the error results, which are displayed as a function of the parameter k (k-NN) and h (Parzen). Indicated on these graphs in the minimum Bayes error estimate produced in leave-one-out testing and the corresponding re-substitution value. The associated error estimates are taken to be the R and L values of error bounding the Bayes error rate.

A.2.6 Sensitivity Analysis.

Currently, the only type of sensitivity analysis that has been automated in the system is that used in exploring the sensitivity of the Bayes error rate to relative errors in the estimate of ambient temperature. The range of relative error and the number of points in this range at which the sensitivity is to be tested must be specified by the user. Additionally, the software allows the user to control the Bayes error estimation by specifying ranges for the parameters k and h. The software then calls the non-parametric Bayes error estimation routines developed by Martin (10) for each point in the range of relative error values.

A.2.7 Quit.

This option is used to escape from the system software, returning the user to the MATLAB input environment.

A.3 Feature Menu and Feature Specialization

Throughout the operation of the detection software, the user is required to specify the feature or features from the feature list to be used for classification. This task is accomplished through the use of a menu that displays the names of each feature. To select a feature, simply press its button on the menu. Once selected, these generalized features must be specialized in terms of wavelength or waveband. Depending upon which of the features is chosen, the user will be required to specify

a set of specialization quantities for each feature. The defaults for these quantities can be selected by pressing 'enter' at each prompt. When providing the system with spectrum sampling numbers, those numbers must be from 1-728. Sample 1 corresponds to 698.1055 wavenumbers and sample 728 corresponds to 3502.0984 wavenumbers, with the samples being evenly spaced in wavenumber.

A.4 Database Searching and Search Specification

Many of the major functions of the detection software require that the JMSP database be searched for spectral observations associated with a particular set of targets and non-targets within a particular set of experiments. In the system, this need is accommodated by the use of files called the target list and the experiment list.

The target list consists of a set of target names, one on each line, exactly as those names will appear in the ground truth files associated with the data. In a separate column on the same line, a target class designator must be listed for each target name. This designator may have either a value of 1 for targets or 0 for non-targets. There is no limit to the number of targets or non-targets that can be contained within a target list.

The experiment list contains two columns with one line per data collection 'experiment' whose spectral observations should be included in the search. An experiment is a set of observations made within a 10-15 minute time interval during the data collection. The first column in the experiment list specifies the data set to which the experiment belongs (e.g. symptom_slew). The second column specifies which experiment in that data set is to be considered in the search (e.g. d06nz5). The database searching routines will expect to find all files associated with an experiment in a directory *dataset/experiment* within the directory specified as the data directory for the system.

To further specify the limitations of the search, the user must also enter the range of observation times which should be included. This is done by entering an earliest and latest hour value into the system when prompted. The default values for earliest and latest hours are 8 and 18, respectively. These hour designations require use of a 24-hour clock. For instance, an hour value of 13 represents 1:00 p.m.

A.5 Detection Report

Upon completion of the classification processing, the results are written to a file referred to as the 'detection report.' A name for this file must be entered upon initiation of the processing or the default name (det_stats.out) must be accepted. The file contains a detailed breakdown of the type of classification performed, features utilized, target classification performance, experiment performance, feature performance and agreement, and summary performance metrics for the set of observations processed. The report also details the number of spectra discovered during the database search for each class. An example of the detection report generated when processing the symptom slew unobscured data set with the normalized spectral radiance wavenumber moment 1 feature is shown in figure 18.

```
                         UNCLASSIFIED

          T A R G E T   D E T E C T I O N   R E P O R T

CLASSIFICATION SCHEME:    Single Feature:  MAP CRITERION

CLASSIFICATION FEATURE LISTING:

(1) Normalized Spectral Radiance Wavenumber Moment 1:  9.097 - 11.08 microns

    Feature PDF file:  slew_clear/norm_spec_rad_mom.pdf


Total run time:  172.3 seconds

Experiment list:  slew_clear/experiment_list


         TARGET LIST AND OBSERVATION STATISTICS

TARG NAME         CLASS  OBS  SPECTRA  TMIN   TMAX    CC   IC
Scud_B_Missile      1      2     8     9.15  11.08     8    0
CARC_Green_Panel    1     11    11     8.30  11.20    11    0
Lance_Missile       1      6     6     9.15  11.08     6    0
UHaul_Truck         1      6    18     8.62  11.08    18    0
M35_Truck           1      8    96     8.62  11.08    96    0
Scud_B_Decoy        1      2     8     9.15  11.08     8    0
M50_Tank            1     13   101     8.62  14.17   100    1
MA2543_Decoy        1     14    94     8.30  14.17    84   10
M752_Launcher       1     14    63     8.30  14.17    62    1
MA2543_TEL          1     24   388     8.20  14.30   365   23
Tree_Canopy         0     18   486     8.30  14.17   458   28


         EXPERIMENT LIST AND DETECTION STATISTICS

EXPERIMENT NAME       TGT  NTGT  CONT     Pe      LOWER   UPPER    CC     Pd      IC     FAR
symptom_slew/a06np5    1     0  0.0000  0.0000  0.0000  0.0000    8   1.0000    0    0.0000
symptom_slew/b06ns4   52     0  0.0000  0.1346  0.0418  0.2274   11   0.8654    0    0.0000
symptom_slew/b06ns5   52     0  0.0000  0.0385  0.0000  0.0907    6   0.9615    0    0.0000
symptom_slew/c06ns4   16    36  0.1585  0.0962  0.0160  0.1763   18   1.0000    0    0.2381
symptom_slew/c06ns5   16    36  0.9467  0.0192  0.0000  0.0566   96   1.0000    0    0.0588
symptom_slew/d06nz4   68    18  0.2704  0.0116  0.0000  0.0343    8   0.9853    0    0.0588
symptom_slew/d06nz5   68    18  0.5400  0.0465  0.0020  0.0910  100   0.9412    1    0.0000
symptom_slew/e06np4    1     0  0.0000  0.0000  0.0000  0.0000   84   1.0000   10    0.8824
symptom_slew/e06np5    1     0  0.0000  0.0000  0.0000  0.0000   62   1.0000    1    0.0000
symptom_slew/f06mz2   18    18  0.0360  0.7917  0.6292  0.9541   365  0.3333   23    0.0000
symptom_slew/f06mz4   18    18  0.4222  0.0000  0.0000  0.0000   458  1.0000   28    0.0000
symptom_slew/f06mz5   18    18  0.4961  0.0000  0.0000  0.0000        1.0000         0.0000
symptom_slew/a07np5    6     0  0.0000  0.0000  0.0000  0.0045        0.9231         0.0000
symptom_slew/b07ns4   52     0  0.0000  0.0769  0.0000  0.1494        1.0000         0.0000
symptom_slew/b07ns5   52     0  0.0000  0.0192  0.0000  0.0566        0.9808         0.0000
symptom_slew/c07ns4   16    36  0.3278  0.0192  0.0000  0.0566        1.0000         0.0588
symptom_slew/c07ns5   16    36  0.4021  0.0192  0.0000  0.0566        1.0000         0.0588
symptom_slew/d07nz4   82    18  0.3172  0.0000  0.0000  0.0000        1.0000         0.0000
symptom_slew/d07nz5   82    18  0.4840  0.0000  0.0000  0.0000        1.0000         0.0000
symptom_slew/e07np4    1     0  0.0000  0.0000  0.0000  0.0000        1.0000         0.0000
symptom_slew/e07np5    1     0  0.0000  0.0000  0.0000  0.0000        1.0000         0.0000
symptom_slew/f07mz4    6    18  0.3010  0.0000  0.0000  0.0000        1.0000         0.0000
symptom_slew/f07mz5   18    18  0.4249  0.0000  0.0000  0.0000        1.0000         0.0000
symptom_slew/h09nza   27     0  0.0000  0.0000  0.0000  0.0000        1.0000         0.0000
symptom_slew/h09nzb   27     0  0.0000  0.0000  0.0000  0.0000        1.0000         0.0000
symptom_slew/h09nzc   27    36  0.2428  0.1489  0.0472  0.2507        0.7273         0.3333
symptom_slew/i09nz9   11    36  0.4223  0.0426  0.1003  0.9091        0.9091         0.0909
symptom_slew/i09nza   11    36  0.4223  0.0000  0.1003  0.9091        0.9091         0.0909
symptom_slew/i09nzb   11    36  0.5666  0.0213  0.0625  1.0000        1.0000         0.0000
symptom_slew/i09nzc   11    36  0.6702  0.0000  0.0625  0.9091        0.9091         0.0000
symptom_slew/i09nzd   11    36  0.6149  0.0213  0.0625  0.9091        0.9091         0.0000
symptom_slew/j09ns1   11     0  0.0000  0.0909  0.2608  0.9091        0.9091         0.0000
symptom_slew/j09ns2   11     0  0.0000  0.4545  0.1603  0.7488        0.5455         0.0000


              OBSERVATION AND DETECTION SUMMARY

Number of Experiments Processed:           107
Number of Target Observations:              96
Number of Target Spectra :                 793
Number of Non-target Observations:          18
Total Number of Target Spectra:            486
Total Number of Non-target Spectra:        114
Total Number of Observations:             1279


              CLASSIFICATION CONFUSION MATRIX

          RESUBSTITUTION                    LEAVE-ONE-OUT

            TARGET   NONTARGET               TARGET   NONTARGET
   TARGET   0.9559    0.0441       TARGET    0.9470    0.0530
NONTARGET   0.0576    0.9424    NONTARGET    0.0638    0.9362


Total Probability of Error:    0.0493    Total Probability of Error:    0.0571
95% Conf. Interval Lower:      0.0374    95% Conf. Interval Lower:      0.0444
95% Conf. Interval Upper:      0.0611    95% Conf. Interval Upper:      0.0698
Probability of Detection:      0.9559    Probability of Detection:      0.9470
      False Alarm Rate:        0.0356          False Alarm Rate:        0.0396


Resubstitution & Leave-one-out agreement metric:   0.9922

Detection statistics stored in file:  slew_clear/norm_spec_rad_mom_stats.out


                         UNCLASSIFIED
```

Figure 18. Example of target detection report generated by detection software.

Appendix B. Symptom Slew Unobscured Test Set

Throughout the presentation of results in Chapter IV, extensive use is made of the symptom slew unobscured test set. Listed in table 15 are the target names and class designations contained in the target list for this data set. Table 16 gives a list of the 107 experiments from which spectral observations were taken. The specific

Target Name	Class
Scud_B_Missile	1
CARC_Green_Panel	1
Lance_Missile	1
UHaul_Truck	1
M35_Truck	1
Scud_B_Decoy	1
M50_Tank	1
MAZ543_Decoy	1
M752_Launcher	1
MAZ543_TEL	1
Tree_Canopy	0

Table 15. Target list utilized in classification processing of symptom slew unobscured test set

set of experiments included in a particular database search depends upon the time limitations specified by the user when the search is initiated. The performance results in Chapter IV, for instance, were generated using just those observations that were made between 8:00 a.m. and 6:00 p.m.

symptom_slew/	symptom_slew/	symptom_slew/	symptom_slew/
a06np1	f06mz3	e07np1	i09nz5
a06np2	f06mz4	e07np2	i09nz6
a06np3	f06mz5	e07np3	i09nz7
a06np4	g06np1	e07np4	i09nz8
a06np5	g06np2	e07np5	i09nz9
b06ns1	g06np3	f07mz1	i09nza
b06ns2	g06np4	f07mz2	i09nzb
b06ns3	a07np1	f07mz3	i09nzc
b06ns4	a07np2	f07mz4	i09nzd
b06ns5	a07np3	f07mz5	j09ns1
c06ns1	a07np4	h09nz1	j09ns2
c06ns2	a07np5	h09nz2	j10ns1
c06ns3	b07ns1	h09nz3	k10nz1
c06ns4	b07ns2	h09nz4	k10nz2
c06ns5	b07ns3	h09nz5	k10nz3
d06nz1	b07ns4	h09nz6	k10nz4
d06nz2	b07ns5	h09nz7	k10nz5
d06nz3	c07ns1	h09nz8	k10nz6
d06nz4	c07ns2	h09nz9	k10nz7
d06nz5	c07ns3	h09nza	k10nz8
e06np1	c07ns4	h09nzb	k10nz9
e06np2	c07ns5	h09nzc	k10nza
e06np3	d07nz1	h09nzd	k10nzb
e06np4	d07nz2	i09nz1	k10nzc
e06np5	d07nz3	i09nz2	k10nzd
f06mz1	d07nz4	i09nz3	k10nze
f06mz2	d07nz5	i09nz4	

Table 16. List of experiments contained in the symptom slew unobscured test set

105

Bibliography

1. Boyd, Robert W. *Radiometry and the Detection of Optical Radiation*. Orlando: Academic Press, 1983.

2. Cederquist, J. N. et al. *Infrared Multispectral Sensor Program, Phase 2: Field Measurements, Analysis and Modeling*. Technical Report WL-TR-94-1095, Wright Laboratory AARI-4, May 1994. Volume 1: Fourier Transform Spectrometer Sensor Characterization.

3. Devijver, Pierre A. and Josef Kittler. *Pattern Recognition: A Statisitical Approach*. London: Prentice-Hall International, 1982.

4. Duda, Richard O. and Peter E. Hart. *Pattern Classification and Scene Analysis*. New York: John Wiley and Sons, 1973.

5. Fix, E. and J. L. Hodges. *Discriminatory Analysis, Nonparametric Discrimination, Consistency Properties*. Technical Report 4, Project 21-49-004, Randolph Field, TX: School of Aviation Medicine, 1951.

6. Fukunaga, Keinosuke. *Introduction to Statistical Pattern Recognition* (Second Edition). San Diego: Academic Press, 1990.

7. Fukunaga, Keinosuke and Donald M. Hummels. "Bayes Error Esimation Using Parzen and k-NN Procedures," *IEEE Transactions of Pattern Analysis and Machine Intelligence*, *PAMI-9*(5) (September 1987).

8. Goodman, Josesph W. *Introduction to Fourier Optics*. New York: McGraw-Hill, 1968.

9. Lachenbruch, Peter A. "An Almost Unbiased Method of Obtaining Confidence Intervals for the Probability of Misclassification in Discriminant Analysis," *Biometrics*, *23* (1967).

10. Martin, Curtis E. *Non-parametric Bayes Error Estimation for UHRR Target Identification*. MS thesis, Air Force Institute of Technology, 1993.

11. Parzen, Emanuel. "On Estimation of a Probability Density Function and the Mode," *Ann. Math. Stat.*, *33* (1962).

12. Peebles, Peyton Z. *Probability, Random Variables, and Random Signal Principles* (Second Edition). New York: McGraw-Hill, 1987.

13. Roggemann, Michael C. *Multiple Sensor Fusion for Detecting Targets in FLIR and Range Images*. PhD dissertation, Air Force Institute of Technology, 1989.

14. Schowengerdt, Robert A. *Techniques for Image Processing and Classification in Remote Sensing*. Orlando: Academic Press, 1983.

15. Smith, Frederick G. "The Infrared and Electro-Optical Systems Handbook." *Atmospheric Propagation of Radiation Volume 2*, edited by Joseph S. Accetta and David L. Shumaker, Ann Arbor: Infrared Information Analysis Center, 1993.

16. Thomas, Jason E. *Multispectral Detection of Ground Targets in Highly Correlated Backgrounds*. MS thesis, Air Force Institute of Technology, 1994.

17. Yakowitz, Sidney and Ferenc Szidarovsky. *An Introduction to Numerical Computations* (Second Edition). New York: Macmillan, 1989.

18. Young, Hugh D. *Statistical Treatment of Experimental Data*. New York: McGraw-Hill, 1962.

19. Yu, Xiaoli, et al. "Comparative Performance Analysis of Adaptive Multispectral Detectors," *IEEE Transactions on Signal Processing, 41*(8) (August 1993).

Vita

Captain Matthew R. Whiteley was born 25 May 1969 in Marietta, Ohio. He graduated from Fort Frye High School as Valedictorian in 1987. Whiteley then attended Carnegie Mellon University in Pittsburgh, Pennsylvania where he received a B.S. in Physics with University Honors. Upon graduation, he accepted a commission as a Second Lieutenant in the United States Air Force. In September, 1991, Whiteley was assigned to the Advanced Guidance Division, Wright Laboratory Armament Directorate, Eglin Air Force Base, Florida. While assigned to the Armament Directorate, he became Chief Engineer of the Advanced Technology Ladar System (ATLAS), a laser radar terminal guidance system. During this time, Whiteley also served as the president of the Eglin Company Grade Officers' Council and the government chairman of the Automatic Target Recognizer Working Group Evaluation Committee. In May, 1994, he was assigned to the Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio to pursue a Master of Science degree in Engineering Physics. Upon completion of that assignment, Captain Whiteley will enter the doctoral program in the Department of Engineering Physics at AFIT.

Permanent address: P.O. Box 6
Whipple, OH 45788

REPORT DOCUMENTATION PAGE

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE September 1995	3. REPORT TYPE AND DATES COVERED Master's Thesis

4. TITLE AND SUBTITLE

Non-Imaging Infrared Spectral Target Detection

5. FUNDING NUMBERS

6. AUTHOR(S)

Matthew R. Whiteley, Capt, USAF

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Air Force Institute of Technology
WPAFB OH 45433-6583

8. PERFORMING ORGANIZATION REPORT NUMBER

AFIT/GAP/ENP/95S-01

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Capt Ray O. Johnson
WL/AARI, WPAFB OH 45433-7408

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

Automatic detection of time-critical mobile targets using spectral-only infrared radiance data is explored. A quantification of the probability of detection, false alarm rate, and total error rate associated with this detection process is provided. A set of classification features is developed for the spectral data, and these features are utilized in a Bayesian classifier. The results of this processing are presented and sensitivity of the class separability to target set, target configuration, diurnal variations, mean contrast, and ambient temperature estimation errors is explored. This work introduces the concept of *atmospheric normalization* of classification features, in which feature values are normalized using an estimate of the ambient temperature surrounding the target. Classification testing of spectral field measurements made on an array of U.S. and foreign military assets reveal a total error rate near 5% with a 95% probability of detection and a concurrent false alarm rate of 4% when a single classification feature is employed. Sensitivity analysis indicates that the probability of detection is reduced to 70-75% in the hours preceding daylight, and that for the total error rate to be less than 10%, the target-to-background mean contrast must be greater than 0.1. Analysis of the atmospheric normalization technique reveals that in order to keep the total error rate less than 10%, the ambient temperature must be estimated with less than 3K absolute accuracy.

14. SUBJECT TERMS

Multispectral, Spectral, Infrared, Target Detection, Critical Mobile Targets, Remote Sensing, Radiometry, Pattern Recognition, Bayesian Classifier

15. NUMBER OF PAGES
120

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

GENERAL INSTRUCTIONS FOR COMPLETING SF 298

The Report Documentation Page (RDP) is used in announcing and cataloging reports. It is important that this information be consistent with the rest of the report, particularly the cover and title page. Instructions for filling in each block of the form follow. It is important to *stay within the lines* to meet *optical scanning requirements*.

Block 1. Agency Use Only *(Leave blank)*.

Block 2. Report Date. Full publication date including day, month, and year, if available (e.g. 1 Jan 88). Must cite at least the year.

Block 3. Type of Report and Dates Covered. State whether report is interim, final, etc. If applicable, enter inclusive report dates (e.g. 10 Jun 87 - 30 Jun 88).

Block 4. Title and Subtitle. A title is taken from the part of the report that provides the most meaningful and complete information. When a report is prepared in more than one volume, repeat the primary title, add volume number, and include subtitle for the specific volume. On classified documents enter the title classification in parentheses.

Block 5. Funding Numbers. To include contract and grant numbers; may include program element number(s), project number(s), task number(s), and work unit number(s). Use the following labels:

C	-	Contract	**PR** -	Project
G	-	Grant	**TA** -	Task
PE	-	Program Element	**WU** -	Work Unit Accession No.

Block 6. Author(s). Name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. If editor or compiler, this should follow the name(s).

Block 7. Performing Organization Name(s) and Address(es). Self-explanatory.

Block 8. Performing Organization Report Number. Enter the unique alphanumeric report number(s) assigned by the organization performing the report.

Block 9. Sponsoring/Monitoring Agency Name(s) and Address(es). Self-explanatory.

Block 10. Sponsoring/Monitoring Agency Report Number. *(If known)*

Block 11. Supplementary Notes. Enter information not included elsewhere such as: Prepared in cooperation with...; Trans. of...; To be published in.... When a report is revised, include a statement whether the new report supersedes or supplements the older report.

Block 12a. Distribution/Availability Statement. Denotes public availability or limitations. Cite any availability to the public. Enter additional limitations or special markings in all capitals (e.g. NOFORN, REL, ITAR).

DOD	-	See DoDD 5230.24, "Distribution Statements on Technical Documents."
DOE	-	See authorities.
NASA	-	See Handbook NHB 2200.2.
NTIS	-	Leave blank.

Block 12b. Distribution Code.

DOD	-	Leave blank.
DOE	-	Enter DOE distribution categories from the Standard Distribution for Unclassified Scientific and Technical Reports.
NASA	-	Leave blank.
NTIS	-	Leave blank.

Block 13. Abstract. Include a brief *(Maximum 200 words)* factual summary of the most significant information contained in the report.

Block 14. Subject Terms. Keywords or phrases identifying major subjects in the report.

Block 15. Number of Pages. Enter the total number of pages.

Block 16. Price Code. Enter appropriate price code *(NTIS only)*.

Blocks 17. - 19. Security Classifications. Self-explanatory. Enter U.S. Security Classification in accordance with U.S. Security Regulations (i.e., UNCLASSIFIED). If form contains classified information, stamp classification on the top and bottom of the page.

Block 20. Limitation of Abstract. This block must be completed to assign a limitation to the abstract. Enter either UL (unlimited) or SAR (same as report). An entry in this block is necessary if the abstract is to be limited. If blank, the abstract is assumed to be unlimited.